

Lagrangian Coherent Data Assimilation for chaotic geophysical systems

Rose Crocker

April 23, 2021

*Thesis submitted for the degree of
Master of Philosophy
in
Applied Mathematics
at The University of Adelaide
Faculty of Engineering, Computer and Mathematical Sciences
School of Mathematical Sciences*



THE UNIVERSITY
of ADELAIDE

Contents

Signed Statement	xiii
Acknowledgements	xv
Dedication	xvii
Abstract	xix
1 Introduction	1
1.1 Making predictions about geophysical systems	1
1.2 Chaos in geophysical models	4
1.3 Thesis overview	7
2 Lagrangian Coherent Structure Theory	11
2.1 Lagrangian Coherent Structures	11
2.2 Defining LCS mathematically	15
2.3 LCS Extraction methods: FTLE	17
2.4 FTLE ridges and LCS	19
2.5 Issues with FTLEs	21
2.6 Numerical calculation of FTLE fields	22
2.7 The FTLE for Rossby Wave Flow	23
2.8 The parameter ϵ and the chaoticity of Rossby wave flow	28
3 Data Assimilation	35
3.1 Data Assimilation for Geophysical Systems	35
3.2 DA Methods	39
3.3 Bayesian Parameter Estimation	42
3.4 Parameter estimation methods	43
3.5 Particle Filters	45
3.6 ABC Algorithms	48
3.7 Lagrangian Data Assimilation	50
3.8 Lagrangian Coherent Data Assimilation	52

4	FTLE-ABC	55
4.1	The FTLE-ABC Algorithm	55
4.2	Research questions for FTLE-ABC	58
4.3	Standard ABC versus FTLE-ABC	60
4.3.1	Spatially stationary structures with uniform initialisation	61
4.3.2	Spatially stationary structures with non-uniform initialisation	64
4.3.3	Spatially traversing structures with uniform initialisation	68
4.4	Standard PF versus FTLE-ABC	69
4.4.1	No initial noise and small versus large time steps	70
4.4.2	Added initial noise and small versus large time steps	76
5	Quantifying FTLE-ABC's performance in chaotic regimes	81
5.1	The root-mean squared error	81
5.2	Comparing RMSEs: standard ABC comparison	82
5.3	RMSE ratios: standard ABC comparison	86
5.4	Comparing RMSEs: standard PF comparison	88
5.5	RMSE ratios: standard PF comparison	91
6	Summary, Discussion and Conclusions	95
6.1	Summary: LaCoDA and FTLE-ABC	95
6.1.1	FTLE-ABC versus standard ABC	96
6.1.2	FTLE-ABC versus standard PF	98
6.2	Further research pathways	99
	Bibliography	103

List of Figures

1.1	Satellite data of the night-time total column carbon monoxide retrieved over Asia. This illustrates the typical gaps in data collected via satellite. These gaps are attributable to both the satellite path as well as blocking due to clouds and other such measurement issues. (From Figure 1 in Lahoz & Schneider (2014))	2
1.2	The progression of resolution from the first IPCC report (FAR, 1990), through the second (SAR, 1996), third (TAR, 2001) and fourth (AR4, 2007) reports. As can be seen, the resolution increases by around a factor of 5 from FAR to AR4, such that the UK is not resolved in the first image but can be clearly seen in the last. (Taken from Figure 1.4 in IPCC AR4 (2007), available at https://www.ipcc.ch/report/ar4/wg1/)	3
1.3	The Lorenz attractor, a famous visualisation of the connection between solution structure and chaotic behaviour.	5
1.4	A comparison of the time evolution of a selection of different ensembles of initial conditions. An initial ensemble is a group of initial points close by on the attractor, denoted by a round red circle. The other red lines show how the initial ensemble is warped as it traverses the attractor. Some initial circles become more warped than others, depending on their initial position on the attractor. The bottom case, for example, becomes very spread out. Some circles also traverse further in the same time, such as comparing the upper left and right cases. Yellow dots are <i>fixed points</i> of the attractor, which determining the shape and direction of trajectories in their neighbourhood. (Derived from Figure 3 in Baines (2008).)	6
2.1	Four examples of Lagrangian Coherent Structures in nature. (a) an observed sudden extension of the Gulf Oil Spill, which was found to possess a coherent <i>strain</i> structure, (b) Jupiter’s Great Red Spot, a coherent rotational structure surrounded by coherent shear structures, (c) a nutrient eddy in the ocean. (d) Hurricane Florence, a very impactful rotational coherent structure. (Images from <i>NASA Image of the day Gallery</i> (2020)).	12

- 2.2 Two common examples of fixed points in steady dynamical systems. The blue shapes represent a fluid parcel and its deformation as it moves with respect to the fixed point. At left is a saddle fixed point, characterised by stretching in one direction and contraction in the orthogonal direction. At right is a centre fixed point, characterised by rotational motion around the fixed point. (Adapted from Figure 2 in Haller (2015)). 14
- 2.3 Schematic of numerically approximating the flow map gradient via gridding. A grid is drawn around each fluid particle's initial position, using deformations in the x direction ($\pm\delta x$) and y direction ($\pm\delta y$). This grid is propagated forward from t_0 to t using the flow map. The original and final grid positions, and the magnitude of δx and δy , are used to approximate the flow map gradient over $[t_0, t]$ 23
- 2.4 (a) shows the phase space diagram for Ψ in (2.8) with $\epsilon = 0$ in Equation 2.7. Saddle-type fixed points exist at P_1, P_2, P_5 and P_6 , tending to cause fluid particle stretching nearby. Centre fixed points are found at P_3 and P_4 , tending to cause fluid particle rotation nearby (Adapted from)Pierrehumbert (1991a)). (b) shows an FTLE field for Ψ in (2.8) with $\epsilon = 0$ for $[t_0, t] = [0, 24]$. It can be seen that FTLE 'ridges' (yellow/orange) draw out boundaries between the stream region, where stretching dominates, and gyre regions, where rotation dominates. 24
- 2.5 FTLE fields for Ψ in (2.8) with $\epsilon = 0$ based on flow times (a) $[0, 3]$, (b) $[0, 6]$, (c) $[0, 12]$ and (d) for $[0, 24]$. From this sequential extension of the calculation time it can be seen how the FTLE ridge (orange/yellow) slowly extends to define the transport barrier between gyres and inner stream. . . . 25
- 2.6 Lagrangian trajectories of 36 fluid particles initially placed on a uniform grid for the time $[t_0, t] = [0, 6]$. Different colours indicate trajectories of different initial positions (indicated with open black circles) on the grid. (a) for the stream function Ψ in the original frame with $\epsilon = 0$. (b) in co-moving frame with $\epsilon = 0$, (c) in the original frame with $\epsilon = 0.6$ and (d) in the co-moving frame with $\epsilon = 0.6$. It can be seen that for $\epsilon = 0$ in the co-moving frame the trajectories mirror the phase space diagram for the steady system shown in Figure 2.4a. 31
- 2.7 FTLE fields for $[t_0, t] = [0, 12]$ for (a) $\epsilon = 0$, (b) $\epsilon = 0.1$, (c) $\epsilon = 0.3$ and (d) $\epsilon = 0.5$. These small incremental increases in ϵ illustrate the sensitivity of the time perturbed Rossby wave model to changes in ϵ . Chaotic behaviour appears to occur more quickly for higher ϵ , a phenomenon discussed in detailed in Section 2.8. 32

2.8 $1/FTLE_{med}$ (orange dashed) and $1/FTLE_{av}$ (blue fill) curves estimate the chaotic timescale for a range of ϵ and for different total calculation times T_f . As can be seen, for $T_f \leq 6$ the chaotic timescale decreases with increasing ϵ , while for $T_f > 6$ the chaotic timescale decreases until it is minimised by a particular ϵ_{min} 33

2.9 Table showing how the FTLE field varies with the calculation time T_f and the parameter ϵ 34

3.1 A schematic of how data assimilation combines information from a mathematical model with collected observational data. Ozone data (an *ozone observation*) is combined with a mathematical model for the geographical distribution of ozone (an *ozone forecast*). Both observations and forecast contain errors associated with data collection methods and model assumptions respectively. The forecast and observations are combined using a data assimilation scheme to form the *ozone analysis*, a prediction for global ozone levels at the time of the forecast and observation. This analysis also has associated error arising from a combination of measurement error, model assumptions and the data assimilation scheme itself. (Adapted from Figure 2 in Lahoz & Schneider (2014)) 37

3.2 Diagram showing the operation of the observation operator $h(\cdot)$ and the model update \mathcal{M}_{n+1} in a generic DA scheme. The observation operator transfers the model state vector at time t_n , \mathbf{x}_n , from the model space, of dimension M , to observation space, of dimension $N \leq M$. Formulation of the analysis update to \mathbf{x}_n is then carried out in observation space. \mathbf{x}_{n+1} is then calculated by evolving the analysis update of the state forward in time using \mathcal{M}_{n+1} . This is repeated at each assimilation step. 38

3.3 Schematic showing one assimilation step of the Kalman filter. A prediction \mathbf{x}_n^p at time t_n is produced by evolving the analysis from the previous time step forward in time using the model \mathcal{M}_n . The quantity \mathbf{x}_n^p is transferred to observation space using the observation operator $h(\cdot)$. The model in observation space $h(\mathbf{x}_n^p)$ is compared with \mathbf{y}_n , data collected at time t_n , via the innovation $I_n = \mathbf{y}_n - h(\mathbf{x}_n^p)$. The analysis \mathbf{x}_n^a is then formulated in model space by acting K_n , the Kalman matrix, on I_n , which weights the elements of I_n while also transferring I_n to model space. The process is repeated for each time step in the assimilation scheme. 41

- 3.4 Diagram of the process of a simple particle filter. To estimate the posterior, various realisations of the posterior are sampled, referred to as ‘particles’. This can be done by sampling θ_n^i from a prior and then using θ_n^i within the model to compute the corresponding \mathbf{x}_n^i . The particles are weighted according to how closely they match \mathbf{y}_t , the observed data. Heavily weighted particles are then duplicated to retain the same number of members in the particle ensemble. Resampling resets the weights to be uniform across the ensemble, so that particle degeneracy is avoided. 45
- 3.5 Schematic of one assimilation step of an Approximate Bayesian Computation (ABC) type DA scheme. Multiple realisations of the parameter θ_n are drawn from the prior distribution and used within the model \mathcal{M}_n to compute forecasts \mathbf{x}_n^i , for each θ_n^i . These forecasts are compared to collected observational data \mathbf{y}_n via a distance measure ρ , such as an L²-norm. Forecasts and observations which most closely match according to the chosen measure are retained and their associated values of θ_n^i used to refine the assumed distribution of θ_n . This is repeated at each time step. 49
- 4.1 Schematic of one assimilation step of the FTLE-ABC DA scheme. Realisations of the parameter ϵ are drawn from a prior distribution (in this case uniform). For each realisation the system is propagated forward in time using the gridding method explained in Figure 2.3. This is used to calculate the flow map gradient and hence an FTLE field for each parameter realisation, which will form the forecast FTLE. An observational FTLE field is also calculated from tracer data collected at the same time. The forecast and observed FTLE fields are compared via an L²-norm distance measure. The best $\alpha\%$ of these ϵ values according to the distance measure are retained to form the prior for the next assimilation step 56
- 4.2 Comparison of the first assimilation step (left) and last assimilation step (right) for FTLE-ABC (blue and orange diamonds) and standard ABC (green and purple crosses). The lines represent the average of the best 20% of particles according to distance for FTLE-ABC (black) and standard ABC (black dashed). ϵ_{true} is denoted with a red cross. (a), (b) use $\epsilon = 0.3$, (c), (d) $\epsilon = 0.9$ and (e), (f) $\epsilon = 1.3$. The runs are for the model in the co-moving frame with $T_f = 12$, $J = 9$, $M = 100$, when initialised on a uniform grid. 61
- 4.3 Comparison of the first assimilation step (left) and last assimilation step (right) for (a), (b) use $\epsilon = 0.3$, (c), (d) $\epsilon = 0.9$ and (e), (f) $\epsilon = 1.3$, with the same colour scheme as Figure 4.2. The runs are for the model in the co-moving frame with $T_f = 12$, $J = 3$, $M = 100$, when initialised on a uniform grid. 62

- 4.4 Comparison of the first assimilation step (left) and last assimilation step (right) for (a), (b) use $\epsilon = 0.3$, (c), (d) $\epsilon = 0.9$ and (e), (f) $\epsilon = 1.3$, with the same colour scheme as Figure 4.2. The runs are for the model in the co-moving frame with $T_f = 12$, $J = 9$, $M = 100$, when initialised inside gyres. 65
- 4.5 Comparison of the first assimilation step (left) and last assimilation step (right) for (a), (b) use $\epsilon = 0.3$, (c), (d) $\epsilon = 0.9$ and (e), (f) $\epsilon = 1.3$, with the same colour scheme as Figure 4.2. The runs are for the model in the co-moving frame with $T_f = 12$, $J = 3$, $M = 100$, when initialised inside gyres. 66
- 4.6 Comparison of the first assimilation step (left) and last assimilation step (right) for (a), (b) use $\epsilon = 0.3$, (c), (d) $\epsilon = 0.9$ and (e), (f) $\epsilon = 1.3$, with the same colour scheme as Figure 4.2. The runs are for the model in the moving frame with $T_f = 12$, $J = 9$, $M = 100$, when initialised on a uniform grid. 67
- 4.7 Comparison of the first assimilation step (left) and last assimilation step (right) for (a), (b) use $\epsilon = 0.3$, (c), (d) $\epsilon = 0.9$ and (e), (f) $\epsilon = 1.3$, with the same colour scheme as Figure 4.2. The runs are for the model in the moving frame with $T_f = 12$, $J = 3$, $M = 100$, when initialised on a uniform grid. 68
- 4.8 Comparison of 6 assimilation steps for the standard PF, in the moving frame with $T_f = 1$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.0$ and $N = 200$. Green crosses are particles, the black dashed line is the particle population average and the red cross is ϵ_{true} 71
- 4.9 Comparison of 6 assimilation steps for FTLE-ABC, in the moving frame with $T_f = 1$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.0$ and $N = 200$. Blue and orange diamonds are particles, with orange representing the best 20% according to distance. The black line denotes the average of the best 20% and the red cross is ϵ_{true} 72
- 4.10 Comparison of 6 assimilation steps for the standard PF, in the moving frame with $T_f = 24$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.0$ and $N = 200$. Green crosses are particles, the black dashed line is the particle population average and the red cross is ϵ_{true} 73
- 4.11 Comparison of 6 assimilation steps for FTLE-ABC, in the moving frame with $T_f = 24$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.0$ and $N = 200$. Blue and orange diamonds are particles, with orange representing the best 20% according to distance. The black line denotes the average of the best 20% and the red cross is ϵ_{true} 74

4.12	Comparison of 6 assimilation steps for the standard particle filter, in the co-moving frame with $T_f = 1$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.02$ and $N = 200$. Green crosses are particles, the black dashed line is the particle population average and the red cross is ϵ_{true}	75
4.13	Comparison of 6 assimilation steps for FTLE-ABC, in the co-moving frame with $T_f = 1$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.02$ and $N = 200$. Blue and orange diamonds are particles, with orange representing the best 20% according to distance. The black line denotes the average of the best 20% and the red cross is ϵ_{true}	76
4.14	Comparison of 6 assimilation steps for the standard particle filter, in the co-moving frame with $T_f = 24$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.02$ and $N = 200$. Green crosses are particles, the black dashed line is the particle population average and the red cross is ϵ_{true}	77
4.15	Comparison of 6 assimilation steps for FTLE-ABC, in the co-moving frame with $T_f = 24$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.02$ and $N = 200$. Blue and orange diamonds are particles, with orange representing the best 20% according to distance. The black line denotes the average of the best 20% and the red cross is ϵ_{true}	78
5.1	Plots of χ_{F-ABC} (orange dashed) and χ_{ABC} (blue filled) over a range of ϵ_{true} values for a selection of ΔT . A fairly consistent peak in RMSE for FTLE-ABC occurs around $\epsilon = 0.5$ and trough around $\epsilon = 0.2$	83
5.2	RMSE contours (a) χ_{F-ABC} and (b) χ_{ABC} against the assimilation time step ΔT and ϵ_{true} . The model is in the moving frame with particles initialised on a standard grid, for a final time of $T_f = 24$. The $1/FTLE_{med}$ vs. ϵ curve for $T_f = 24$ is also plotted in red in (a).	85
5.3	Contours of $\tau_{F/ABC}$ for the model in the moving frame, for a final time of $T_f = 24$, initialised on a uniform grid with $M = 100$ tracers. The red line is the $1/FTLE_{med}$ vs ϵ plot for $T_f = 24$	87
5.4	Plots of χ_{F-ABC} (orange dashed) and χ_{PF} (blue filled) over a range of ϵ_{true} values. As can be seen, for $\Delta T \in [2, 8]$ and $\epsilon_{true} > 0.4$, FTLE-ABC tends to out-perform the standard PF, while for lower ΔT and ϵ the PF performs better.	89
5.5	Contours of RMSE for (a) χ_{F-ABC} and (b) χ_{PF} against the assimilation time step ΔT and ϵ_{true} , for the model in the moving frame with particles initialised on a standard grid, for a final time of $T_f = 24$	90
5.6	Contours of $\tau_{F/PF}$ for the model in the moving frame, for a final time of $T_f = 24$, initialised on a standard grid for $M = 36$ tracers. The red line is the $1/FTLE_{med}$ vs ϵ plot for $T_f = 24$	92

6.1 Table showing how the FTLE field varies with the calculation time $\Delta T = T_f - 0$ and the parameter ϵ 100

Signed Statement

I certify that this work contains no material which has been accepted for the award of any other degree or diploma in my name in any university or other tertiary institution and, to the best of my knowledge and belief, contains no material previously published or written by another person, except where due reference has been made in the text. In addition, I certify that no part of this work will, in the future, be used in a submission in my name for any other degree or diploma in any university or other tertiary institution without the prior approval of the University of Adelaide and where applicable, any partner institution responsible for the joint award of this degree.

I give consent to this copy of my thesis, when deposited in the University Library, being made available for loan and photocopying, subject to the provisions of the Copyright Act 1968.

I also give permission for the digital version of my thesis to be made available on the web, via the University's digital research repository, the Library Search and also through web search engines, unless permission has been granted by the University to restrict access for a period of time.

Signed: Date:

Acknowledgements

Thank you to my supervisors Sanjeeva, Jack and Lewis, for teaching me so much and making research fun. Thanks also to the wonderful people in the School of Mathematical Sciences, students and staff, who made such a supportive environment.

Finally, thank you to friends and family who distracted me when research was going badly and celebrated with me when it was going well.

Dedication

For my parents, who were surprised the child of an artist and a historian wanted a Masters in maths but supported me anyway.

Abstract

This thesis develops a new method for estimating geophysical parameters based on Lagrangian Coherent Data Assimilation (LaCoDA), a nascent field combining data assimilation and Lagrangian coherent structure techniques. Lagrangian coherent structure theory deals with characterising and extracting fluid structures which have a dominant impact on the transport of key flow properties (Balasuriya et. al., *Physica D*, 2018:31-51). Data assimilation (DA) is a methodology for combining information from observational data with that from a mathematical model to make predictions about dynamical systems (Lahoz & Schneider, *Front. Environ. Sci.*, Springer-Verlag, 2014). Lagrangian Coherent Data Assimilation attempts to combine these two areas to devise data assimilation schemes which exploit information from Lagrangian coherent structures to improve data assimilation in chaotic systems (Maclean et. al. *Physica D*, 2017:36-45).

The LaCoDA technique developed here combines the data assimilation algorithm known as Approximate Bayesian Computation (ABC) with a Lagrangian coherent structure method, the Finite Time Lyapunov Exponent (FTLE). The new method, denoted FTLE-ABC, is tested on estimating the ϵ parameter from the Rossby wave model, a mathematical model which simulates an important type of atmospheric flow. FTLE-ABC is shown to outperform the benchmark methods, a standard particle filter and a standard ABC scheme, for particular regimes of the true value of ϵ , the chaoticity of the flow and the time step used in the DA scheme. In particular, the estimated chaotic timescale is found to impact FTLE-ABC's performance, with the algorithm often performing better in parameter regimes for which the chaotic timescale is fairly constant with ϵ .

Chapter 1

Introduction

Making predictions about geophysical systems has crucial applications in governance, economics, agriculture and public health. Modeling geophysical systems is challenging, however, due to erratically distributed data and complex, high-dimensional models. Additionally, geophysical systems exhibit chaotic behaviour which severely impacts prediction accuracy and reliability. Chapter 1 outlines how understanding the mathematics of chaos may allow insight into the predictability of geophysical systems. This motivates the main aim of this thesis, the development of an algorithm which exploits the mathematical structure of time-dependent chaotic systems to improve geophysical predictions.

1.1 Making predictions about geophysical systems

Geophysical systems, such as the oceans, atmosphere and climate, have a significant impact on governance and policy making. The Intergovernmental Panel on Climate Change (IPCC) Reports, which are produced annually to provide recommendations to governments, policy-makers and industry on future climate scenarios, are a well known example of this significance (IPCC 2020). In particular, the agricultural, public health and economic sectors strongly rely on knowledge of geophysical forecasts to produce judicious recommendations and decisions. Making timely and accurate predictions about geophysical systems will only become of increasing concern as issues of climate security and adaptation intensify.

Systems such as the climate, however, continue to present significant challenges to mathematical modelers. One issue is that data on these systems may be erratically distributed geographically, such that parts of the globe are abundant in information while others are sparse. This is illustrated in Figure 1.1, which shows night-time total column carbon monoxide data collected via satellite across Asia. As can be seen, due to the

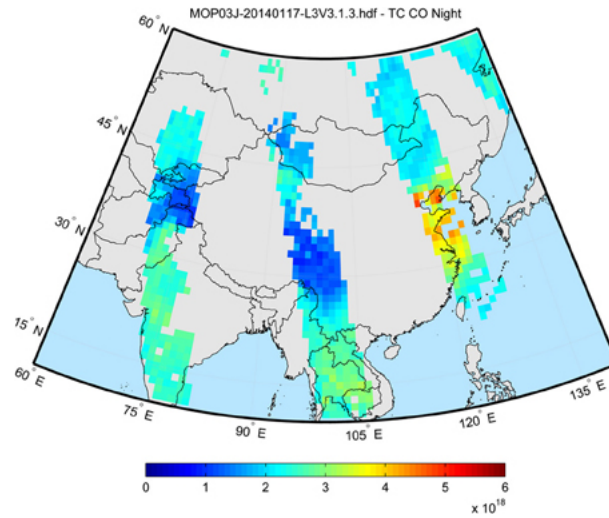


Figure 1.1: Satellite data of the night-time total column carbon monoxide retrieved over Asia. This illustrates the typical gaps in data collected via satellite. These gaps are attributable to both the satellite path as well as blocking due to clouds and other such measurement issues. (From Figure 1 in Lahoz & Schneider (2014))

path of the satellite and cloud coverage, there are large latitudinal strips where data is absent. One common issue modelers face is how to make suitably accurate predictions with limited information from data (Lahoz & Schneider 2014, Kalnay 2002).

A second, crucial issue is that geophysical models and data are often extremely high dimensional, such that running models may be unreasonably computationally expensive (Kalnay 2002, Carrassi et al. 2018). An example of this is shown in Figure 1.2, which depicts the progression of resolution in climate models from the first IPCC assessment report in 1990 (FAR) to the fourth assessment report in 2007 (AR4). The resolution of the models used by the IPCC became around 5 times finer over the 17 years between FAR and AR4, such that the UK is not resolved in the first image, but quite clearly resolved in the last. This significant rate of increase in resolution strongly impacts the computational expense of model running and analysis. Consequently, researchers must continually develop more efficient methods of modeling geophysical systems to keep up with burgeoning increases in dimension.

Lastly, geophysical systems exhibit chaotic properties which strongly impact prediction error. Chaotic behaviour causes model predictions to be highly sensitive to slight errors in the initial model inputs, such that small initial errors are magnified egregiously as the model is propagated through time. The issues of high-dimensionality and sparse data are also often magnified by chaotic behaviour, as chaos severely impacts system




Figure 1.2: The progression of resolution from the first IPCC report (FAR, 1990), through the second (SAR, 1996), third (TAR, 2001) and fourth (AR4, 2007) reports. As can be seen, the resolution increases by around a factor of 5 from FAR to AR4, such that the UK is not resolved in the first image but can be clearly seen in the last. (Taken from Figure 1.4 in IPCC AR4 (2007), available at <https://www.ipcc.ch/report/ar4/wg1/>)

predictability. Increases in model accuracy due to increasing model resolution, for example, are constrained by chaotic behaviour, such that the additional computational expense of increasing resolution may not reap significant gains in accuracy if the system is highly chaotic. The difficulty of making predictions about systems where data is sparse is also worsened in highly chaotic systems, as small data sets will already have significant initial error which is then magnified by chaotic behaviour. Chaos is an innate feature of geophysical systems, however, and improving methods of modeling chaotic systems is foundational to improving geophysical modeling (Sivakumar 2004, Mihailovic et al. 2014).

1.2 Chaos in geophysical models

Chaos is a much studied property of dynamical systems, associated with extreme sensitivity to initial conditions and seemingly random and disordered behaviour. Such perceived randomness is an innate property of many deterministic physical systems and cannot be eliminated by simply improving the model resolution (Alligood et al. 1997). One of the first researchers to identify chaotic behaviour in geophysical systems was the meteorologist Edward Lorenz (Fowler & McGuinness 2019). As the story goes, Lorenz was running computer simulations of a simple model of the atmosphere and began feeding his simulation readouts from previous runs as initial conditions. His readouts, however, were of lower resolution than the computer memory storage, so that the input figures were rounded versions of the previous output. He found that the system behaved in a completely different fashion when previous outputs of the same system were used as inputs, despite these inputs only differing from the originals by three decimal places (Gleick 2008). Based on his studies of simple chaotic models, Lorenz correctly postulated that long-range weather forecasting would be inherently impeded by the non-periodic properties of the atmosphere and related systems. Indeed, a characteristic of chaotic systems is complete unpredictability on sufficiently long timescales (Lorenz 1963) .

Although giving rise to randomness and unpredictability, chaotic behaviour is also deeply connected to structure: the geometric structure of solutions to the differential equations governing a dynamical system. Lorenz's discovery of this feature of chaos was documented in a 1963 paper on a now widely used, yet very simple, model of atmospheric convection (Lorenz 1963). Lorenz plotted the model's spatial solutions x, y, z in three dimensional space as they evolved through time. This generated what is called a strange attractor, shown in Figure 1.3, a complex yet highly structured object consisting of solution trajectories which may orbit infinitesimally close to each other yet never intersect. Attractors may contain *fixed points* or *periodic orbits*, discussed in detail in Chapter 2, which shape the geometry and direction of trajectories in a local neighbourhood. Attractors assist in explaining how a completely deterministic system can exhibit




Figure 1.3: The Lorenz attractor, a famous visualisation of the connection between solution structure and chaotic behaviour.

the ‘extreme sensitivity to initial conditions’ characteristic of chaos. Each trajectory represents a unique, deterministic time-evolution route for the model, yet, given the closeness of such trajectories, a small perturbation to an initial condition may result in a completely different traversal of the solution space. This deeply nested structure of attractors produces the erratic, seemingly random behaviour of chaotic systems (Stewart 1997).

Another characteristic of chaotic systems which may be elucidated from solution structure is how some initial conditions give rise to a greater level of unpredictability than others, which is why the reliability of weather forecasts varies so greatly. The Lorenz attractor can again be used to illustrate this property. Figure 1.4 shows three different initial *ensembles*, or selections of nearby initial conditions, indicated by a round red circle. As each of the ensembles moves along its trajectory on the attractor (indicated by perturbed red circles), some initial selections of conditions, such as the top right scenario, are quickly spread out as they move along the attractor, while others remain reasonably unperturbed, as in the top left scenario. Additionally, the cases at the top left and bottom show very similar initial ensembles, yet the two evolve along very different paths, with the bottom ensemble becoming spread between two wings of the attractor, while the top left ensemble only moves to the right wing. Hence, some initial ensembles are very sensitive to small perturbations, while others are not so much. In this way, analysing the structure of objects like attractors reveals significant information about the predictability of a geophysical system (Baines 2008).

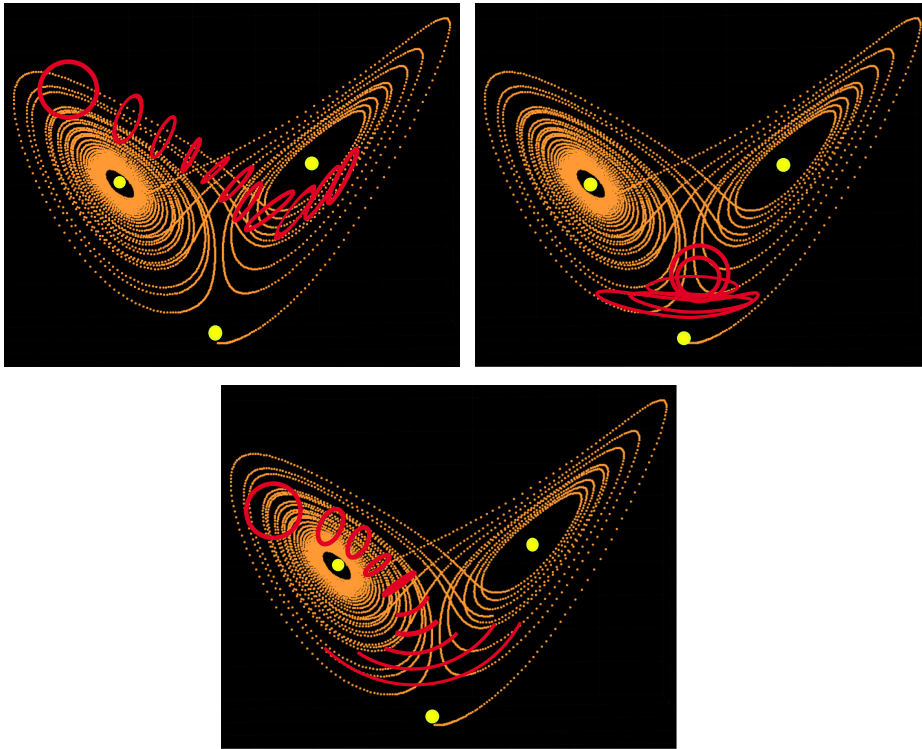


Figure 1.4: A comparison of the time evolution of a selection of different ensembles of initial conditions. An initial ensemble is a group of initial points close by on the attractor, denoted by a round red circle. The other red lines show how the initial ensemble is warped as it traverses the attractor. Some initial circles become more warped than others, depending on their initial position on the attractor. The bottom case, for example, becomes very spread out. Some circles also traverse further in the same time, such as comparing the upper left and right cases. Yellow dots are *fixed points* of the attractor, which determining the shape and direction of trajectories in their neighbourhood. (Derived from Figure 3 in Baines (2008).)

A key question which arises when considering the information captured by solution structure in chaotic systems is *Can we use this structure to quantify unpredictability?* and further, *Can we use quantifications of unpredictability to improve model predictions?* An early development in quantifying uncertainty in chaotic systems, before chaos was recognised as a discipline, was the Lyapunov exponent, originally devised in Alexandr Lyapunov's doctoral thesis (Lyapunov 1892). Lyapunov's work considered regimes of stability and instability in differential equations, rigorously showing that stable solutions exhibit resilience to sufficiently small perturbations, while unstable solutions do not (Bar-

reira 2017). To quantify the stability of a solution, he developed the Lyapunov exponent, a method still used to analyse a variety of dynamical systems. The maximal Lyapunov exponent is effectively the maximal rate of divergence, in the limit as time approaches infinity, between two points which are initially infinitesimally close. This quantifies the expected rate of divergence for an initial condition when a small perturbation is applied. A chaotic system must contain at least one positive maximal Lyapunov exponent, and the magnitude of this suggests the degree of unpredictability in the corresponding initial condition (Alligood et al. 1997).

1.3 Thesis overview

Many established definitions of chaos tend to be in the limit as time approaches infinity. Attractors represent regions of a dynamical system's solution space towards which solutions tend to settle over long times (Stewart 1997). Similarly, Lyapunov exponents represent the maximal divergence rate of initially very close points as time approaches infinity (Barreira 2017). In both cases, transient fluctuations over short times are ignored so as to focus on a system's long-term structure. In practice, time is usually extended to some sufficiently large number so as to reasonably apply infinite-time chaotic theory. This approach, however, is problematic in systems whose governing equation is time-dependent, which is the norm in geophysical systems. In this case an attractor represents a single snapshot of the system structure at a particular time, but gives no information about how such snapshots evolve into each other over time. A mathematical methodology known as *Lagrangian Coherent Structure Theory*, the topic of Chapter 2, has been developed to investigate structures arising in time-dependent systems over *finite times*. Such structures are common in nature, including hurricanes and ocean gyres, and have been shown to contain significant information about geophysical flow behaviour (Balasuriya et al. 2018, Haller 2015). Lagrangian coherent structures can be thought of as finite-time analogs of infinite-time chaotic structures, revealing short-term information about a dynamical system's solution structure. Given the rich information about the long-term behaviour of dynamical systems revealed by infinite-time chaotic structures, analogously, this thesis explores the extent to which Lagrangian Coherent Structure Theory can be used to inform predictions about time-dependent geophysical systems.

A common methodology for making predictions about time-dependent systems is known as *data assimilation*, the topic of Chapter 3. Data assimilation employs mathematics from a range of sub-disciplines, including optimisation, statistics and dynamical systems theory, to judiciously combine information from observational data and a mathematical model (Law et al. 2015). It is the dominant methodology for both predicting the state of a geophysical system and estimating parameters associated with geophysical

systems. A recent subset of the data assimilation literature develops data assimilation methods which may incorporate *Lagrangian data*: data which tracks the individual trajectories of initial conditions in a geophysical flow. Lagrangian data assimilation is of interest as Lagrangian data is becoming increasingly common in areas such as oceanography, and contains rich information regarding extant Lagrangian coherent structures. Current techniques, however, tend to deal badly with highly chaotic, high dimensional and/or highly non-linear models, all of which are common in geophysical systems. Additionally, Lagrangian data assimilation methods which actively exploit the structural information contained in Lagrangian data have not been widely researched. Recent studies, however, have suggested a significant advantage for such methods in chaotic systems and systems with fine scale chaotic structures (Maclean et al. 2017, Gaultier et al. 2013, 2014, Schlueter-Kuck & Dabiri 2019). This thesis builds on a small but rapidly developing literature which attempts to exploit Lagrangian coherent structure theory to improve the performance of data assimilation schemes in chaotic systems. The latter half of Chapter 3 explores this literature to build groundwork for the premise of this thesis: the development of a Lagrangian coherent structure based data assimilation algorithm to enhance parameter estimation in chaotic geophysical models.

There are a variety of numerical methods for extracting Lagrangian Coherent Structures, which are discussed in Chapter 2. Of these methods, a finite-time analog of the Lyapunov exponent, called the Finite Time Lyapunov Exponent (FTLE) (Shadden et al. 2005), is chosen as a basis for the geophysical parameter estimation algorithm developed here. The FTLE has been used as a measure of uncertainty in data assimilation schemes previously (Gaultier et al. 2013, 2014), but not for parameter estimation purposes. Coherent structure methods have also previously been used for geophysical parameter estimation (Maclean et al. 2017, Schlueter-Kuck & Dabiri 2019), but using quite simplistic methods of coherent structure detection which cannot track spatially traversing structures. The FTLE has the advantage of being a frame-independent method, allowing its use in flows with spatially traversing structures. This method is integrated within a common statistical method for data assimilation, called Approximate Bayesian Computation (ABC), to form a new algorithm for parameter estimation which is able to utilise information from a variety of time-dependent chaotic flow structures.

In Chapters 4 and 5 the performance of this new algorithm is tested against two benchmark data assimilation algorithms for parameter estimation, a particle filter and a standard ABC scheme, with a focus on the algorithm's aptitude for prediction when the model is highly chaotic. The model chosen to test the method is the Rossby wave model, introduced in Section 2.7, which is a simple model of an important type of atmospheric flow. As is discussed, this model is useful because the flow demonstrates a variety of spatially traversing structures, and the degree of chaotic behaviour the model exhibits

may be tuned via an input parameter (Pierrehumbert 1991*a*). In particular, Chapters 4 and 5 seek to address the following research questions:

- Does an FTLE-based ABC scheme for parameter estimation offer a performance advantage over the benchmarks in chaotic regimes of Rossby wave flow?
- What is the impact of the time-step length on the performance of the FTLE-based ABC scheme?
- Can the FTLE-based ABC scheme deal with spatially traversing coherent structures, such as spatially moving gyres?
- What is the impact of the initialisation scheme on the performance of the FTLE-based ABC scheme?

Chapter 4 presents an initial investigation into parameter regimes in which the new algorithm out performs the benchmarks, exploring the effect of the initialisation scheme, magnitude of chaoticity, time step length and data noise. Chapter 5 further investigates scenarios of interest identified in Chapter 4, using the statistical measure of root mean squared error to quantitatively compare the new algorithm's performance to the benchmarks. Chapter 6 summarises the results of this thesis, discusses interpretations in the context of Lagrangian coherent structure theory and presents potential pathways for further research.

Chapter 2

Lagrangian Coherent Structure Theory

Lagrangian Coherent Structure (LCS) Theory is a mathematical framework for conceptualising and numerically extracting transient yet impactful structures in time-dependent geophysical systems. Chapter 2 examines why such structures may be seen as finite-time analogues of infinite-time chaotic structures, as well as crucial distinctions in the time-dependent context. Definitions of LCS, applications of coherent structure theory and common techniques used to numerically extract LCS are discussed to provide a theoretical framework for understanding how LCS theory is being used to improve modeling of chaotic geophysical systems. The model of choice in this investigation, Rossby wave flow, is also introduced, with an emphasis on how LCS methods may be used to analyse the flow's chaotic structures.

2.1 Lagrangian Coherent Structures

Lagrangian Coherent Structures (LCS) are large-scale fluid structures that are ubiquitous in nature. Figure 2.1 shows four well-studied structures which may be classified within the LCS framework, including a structure formed during the Deepwater Horizon Oil Spill (Figure 2.1a), Jupiter's Great Red Spot (Figure 2.1b), nutrient eddies in the ocean (Figure 2.1c) and Hurricane Florence (Figure 2.1d). From these images the impact such structures may have on both geophysical systems and associated ecological and human systems can be appreciated. Such structures are often responsible for the transport of key properties in geophysical flows, such as heat, nutrients and energy, and often are a strong determinant of gross system behaviour (Balasuriya et al. 2018). Intuitively, the role of LCS in determining the mixing and transport properties of geophysical flows suggests these structures carry mathematically utilizable information for analysing flow behaviour.




Figure 2.1: Four examples of Lagrangian Coherent Structures in nature. (a) an observed sudden extension of the Gulf Oil Spill, which was found to possess a coherent *strain* structure, (b) Jupiter's Great Red Spot, a coherent rotational structure surrounded by coherent shear structures, (c) a nutrient eddy in the ocean. (d) Hurricane Florence, a very impactful rotational coherent structure. (Images from *NASA Image of the day Gallery* (2020)).

'Coherent structures' have been of interest in fluid mechanics for a long time, both for flow visualisation and as a means of deriving reduced-order models (Balasuriya et al. 2018). Traditionally, an *Eulerian* approach is taken, wherein only time snapshots of fluid particle positions are considered. Generally, this involves analysing the streamlines of the flow, which give the instantaneous tangent to the flow velocity. For *steady* flows, those with a governing equation of the form $\frac{d\mathbf{x}}{dt} = \nu(\mathbf{x})$, we may use classic dynamical systems methods to analyse the flow structure. In this methodology we plot a phase portrait, a grid of spatial coordinates with vectors indicating the direction of the velocity field for the flow at each grid point. Points $\tilde{\mathbf{x}}$ at which $\left. \frac{d\mathbf{x}}{dt} \right|_{\mathbf{x}=\tilde{\mathbf{x}}} = 0$, are called the *fixed points* or

equilibria of the system. The nature of the phase space trajectories around each fixed point determines its categorisation, which in turn may be used to describe the nature of the flow locally. Depending on the type of equilibrium, the phase portrait may include associated *stable* and *unstable manifolds*: material curves on surfaces along which fluid particles asymptote towards fixed points in forwards or backwards time. These manifolds demarcate fluid barriers which separate structurally distinct regions of fluid, and tend to cause nearby fluid parcels to be stretched or contracted. Equilibria may also be enclosed by periodic orbits or families of nested KAM tori, around which fluid particles will tend to rotate, representing rotational, eddy-like structures in the flow (Beron-Vera et al. 2010).

As an illustration of this classic *Eulerian* analysis, the phase portraits for two very common fixed point categorisations are shown in Figure 2.2, a *saddle point equilibrium* (at left) and a *centre fixed point* (at right). In the case of the saddle fixed point, the stable and unstable manifolds form a cross with its centre at the equilibrium. An infinitesimal parcel of fluid, or fluid particle, will tend to travel towards the equilibrium along the stable manifold and away from the equilibrium along the unstable manifold, so that it is stretched out over time. For the centre fixed point, fluid particles tend to orbit the equilibrium while never approaching it, generating rotational flow structures. The possible phase space trajectories become much more complicated with an added third dimension, giving rise to the strange attractor type structures mentioned in Chapter 1, yet such structures still contain *invariant sets* which control how an initial condition evolves through the phase space. A key property of these steady systems is that, although trajectories may be highly complex as in the Lorenz attractor, the phase space trajectories give information about how fluid particles will travel for all times in the flow (Nolte 2019). Consequently, identifying the nature of extant manifolds corresponding to invariant sets allows general knowledge of how fluid particles will behave as they travel within the flow.

However, in unsteady systems, those with a governing equation which is time dependent, the interpretation of equilibria and phase portraits becomes considerably more complex. In this case phase portraits may be thought of as time snapshots of the system, based only on current fluid particle positions and prevailing velocities, without any relation to the system's long-term evolution in time. With time dependence, 'fixed points' of an Eulerian snapshot are no longer temporally fixed, and so the Eulerian approach may give misleading information about the flow structure. Thus, the rich information about flow behaviour discussed in the context of the Lorenz attractor in Chapter 1 is no longer a reliable reference in time-dependent systems. This issue was recognised in a series of investigations by Haller and collaborators (Haller & Yuan 2000, Haller 2002, 2001), who identified key cases where classic phase space analysis failed to give reliable information about flow structure, such as where flows possessed spatially traversing structures.

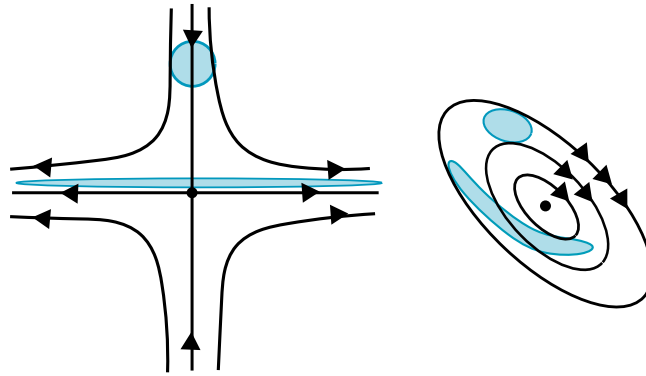


Figure 2.2: Two common examples of fixed points in steady dynamical systems. The blue shapes represent a fluid parcel and its deformation as it moves with respect to the fixed point. At left is a saddle fixed point, characterised by stretching in one direction and contraction in the orthogonal direction. At right is a centre fixed point, characterised by rotational motion around the fixed point. (Adapted from Figure 2 in Haller (2015)).

The key difference between steady and unsteady systems that Haller and his collaborators discovered can be intuited by drawing analogies between classic steady structures and those observed in unsteady systems. In Figure 2.1b, for example, it can be appreciated that Jupiter’s Great Red Spot is analogous to a centre fixed point, in that it possesses a rotationally-dominated structure. Similarly, it can be seen from Figure 2.1a that the LCS observed in the Gulf Oil Spill is analogous to a saddle fixed point, dominated by stretching structures. However, as Haller and his associates point out, entities such as nested streamlines or stable and unstable manifolds identified in snapshots of these systems do not necessarily correspond to unsteady vortices or strain structures. This is because fluid particles in the unsteady scenario are not confined to follow the streamlines of the instantaneous Eulerian snapshot for all times. The velocity at a particular point in space will be changing with time, and so how fluid particles move depends not only on where they are, but the time at which they are at that point. Essentially, in the unsteady case we are interested in the *Lagrangian* perspective; a perspective which incorporates fluid particle trajectories that are solutions to the system’s governing equation. In this framework we keep track of the paths of fluid particles as they are moved through the flow over time, thus incorporating time-dependence into the analysis.

Theory which uses the Lagrangian approach to extract fluid structures from unsteady flows is now known as *Lagrangian Coherent Structure Theory*, a term coined by Haller and Yuan (2000). Rather than assessing fixed points and their associated phase space structures, this alternative approach seeks fluid structures which are “transient and time

evolving yet exhibit some ‘coherence’ over a finite time interval” (Haller 2015). Such an analysis strives to isolate “special surfaces of fluid trajectories that organize the rest of the flow into ordered patterns” (Balasuriya et al. 2018), thus pursuing time-evolving structure in the flow. As such structures require the consideration of a set of fluid trajectories, rather than just initial points, the resulting theory becomes considerably more complex than in the Eulerian perspective. Indeed, numerical methods are commonly used to identify LCS in all but the simplest cases.

Since Haller and associates’ (Haller & Yuan 2000, Haller 2002, 2001) initial developments in the theory of LCS, many researchers from various perspectives have focused on the problem of defining, categorising and numerically extracting LCS. Notable theoretical developments include Lekien and Ross’ (2010) generalisation of LCS to arbitrary Riemann manifolds, allowing the study of LCS on challenging surfaces such as large-scale transport in the curved atmosphere. Many studies have also applied the theory of transfer operators to seek almost invariant manifolds in a range of geophysical flows, combining the theory of Markov chains from stochastic modeling to identifying LCS (Froyland et al. 2014, for example). Due to the area’s dependence on non-analytic methods, studies such as Ross et al. (2010) and Olcay et al. (2010) have investigated the sensitivity of common extraction techniques to data noise, developing methods to overcome associated resolution issues. Many improvements have also been achieved in the computational efficiency of methods. Brunton et al. (2010) devise schemes to significantly increase the efficiency of popular methods, while Lipinski et al. (2010) have developed structure-tracking algorithms that avoid wasting computations on regions too far from key structures. Responding to the field’s increasing applications in science and engineering, Onu, Hahn and Haller (2014) created an LCS toolbox which allows the simple implementation of common LCS extraction algorithms for a variety of applications. As the field continues to develop, numerically extracted LCS are now routinely used in the study of time-dependent fluid systems in a diverse range of areas, including to understand blood transport in the heart (May-Newman et al. 2016), multi-body problems in astrodynamics (Lin et al. 2017) and the development of oceanic algal blooms (Olascoaga et al. 2008).

2.2 Defining LCS mathematically

Unlike fixed points and their associated flow structures in steady systems, there is no universally accepted definition for an LCS. Balasuriya et al. (2018), however, set out a generalised framework for defining LCS which provides a working basis for their mathematical characterisation. The flow is considered within the flow domain $U(t) \in \Omega$ at time t , for Ω the model state space. To implement the framework using Lagrangian machinery, an LCS must be defined at an initial time t_0 , based on the flow from times t_0 to t . Con-

sequently, a crucial object in the Lagrangian framework is the deterministic Lagrangian flow map from times t_0 to t , defined as,

$$\mathcal{F}_{t_0}^t(\mathbf{x}_0) := \mathbf{x}(t_0; t, \mathbf{x}_0) \quad (2.1)$$

with respect to solutions $\mathbf{x}(t_0; t; \mathbf{x}_0)$ to the governing equation

$$\frac{d\mathbf{x}}{dt} = \nu(\mathbf{x}, t), \quad (2.2)$$

with initial condition $\mathbf{x}(t_0; t_0; \mathbf{x}_0) = \mathbf{x}_0$. The flow map takes the set of initial particle positions $\mathbf{x}_0 \in U(t_0)$, and evolves these positions to time t according to Equation (2.2). The flow map allows tracking of fluid particle trajectories over time, making it central to any Lagrangian formulation. Additionally, Balasuriya et al. (2018) require that the fluid velocity field have continuous spatial and temporal derivatives and that all fluid particle trajectories are defined and unique for all initial conditions in $U(t_0)$.

Accepting these requisites, an LCS may be conceptualised as a subset of $U(t_0)$, extracted by examining $\mathcal{F}_{t_0}^t(U(t_0))$. An LCS is therefore dependent not just on the initial time t_0 , but on the time interval over which we consider the flow's evolution. The LCS itself may be identified according to some criterion imposed on the behaviour of the flow map $\mathcal{F}_{t_0}^t(\mathbf{x}_0)$. For example, an LCS could be defined as some subset $S \in \Omega$ such that $\mathcal{F}_{t_0}^t(S)$ is sufficiently similar to S as defined by connectivity and minimal filamentation of the boundary ∂S in the time interval considered. Such a criterion would select for structures like vortex cores, which often possess a coherent core surrounded by a filamenting boundary surface. Another example is seeking S such that $\mathcal{F}_{t_0}^t(S)$ is more strongly attracting than other sets of particles in the neighbourhood of S . This criterion would seek to divide the flow domain into regions with barriers of maximal stretching dividing each region. Essentially, an LCS theory should define a criterion based on the structure it seeks to extract. It follows that, abstractly, an LCS at time t may be defined as a structure resulting from the application of a criterion via some procedure P ,

$$\mathbf{LCS}(t_0; t) := P(\mathcal{F}_{t_0}^t(\cdot)),$$

where $\mathbf{LCS}(t_0; t)$ is a subset of Ω .

The selection of specific criteria and computational procedures capable of applying such criteria to the flow map has led to the development of many LCS extraction methods. There are, for example, trajectory average methods, which compute the time-average of a quantity of interest over Lagrangian trajectories (Balasuriya et al. 2018). This allows a field of Lagrangian averages to be plotted with respect to their corresponding initial conditions, such that exceptionally large or small values in the field suggest extreme values

of the chosen flow quantity and hence a notable structural difference in the flow. Common choices for flow quantities include the norm of the fluid particle speed or velocity. A more complicated trajectory average method, called the Lagrangian Averaged Vorticity Deviation (LAVD), considers the time-averaged norm difference in a fluid particle’s vorticity and the average vorticity over the flow domain. LAVD is particularly useful for elucidating rotationally-dominated structures in a flow, such as gyres and vortices, as its formulation was derived specifically to give a frame-independent quantification of rotation (Haller et al. 2016). Another common approach uses variational methods to define hyperbolic LCS as the most repelling or attracting curves (or surfaces) advected over a finite time in a locality of the flow (Haller 2011, Haller & Beron-Vera 2012). Such surfaces are often derived in terms of the *Cauchy-Green strain tensor*, defined in Section 2.3, which quantifies the strain undergone by particles as they are advected by the flow. Methods using the theory of transfer operators are popular when investigating structures arising in density flows, such as chemical pollutants or salinity (Froyland et al. 2014). These studies approach LCS from a probabilistic standpoint, considering the probability of density transitions between defined ‘boxes’ in the flow and seeking ‘almost invariant’ regions by considering the corresponding transition matrices’ eigenvalues. Another notable method is the recently introduced Stochastic Sensitivity, which quantifies the sensitivity of the flow to stochastic noise. This approach elucidates coherent structures as ridges of maximal sensitivity of Lagrangian trajectories to small perturbations. Stochastic Sensitivity is of particular interest in the context of oceanographic and atmospheric models as these often rely on noisy Lagrangian velocity data and hence cannot assume perfect knowledge of Lagrangian trajectories (Balasuriya 2020a). Perhaps the most commonly used method in geophysical applications, however, is the Finite Time Lyapunov Exponent (FTLE), which measures the magnitude of stretching fluid particles undergo as they are advected by the flow. Due to its wide and varied use in the coherent structure literature the FTLE is the method of choice in this investigation.

2.3 LCS Extraction methods: FTLE

As mentioned, the FTLE method uses a criterion which seeks structures with an exceptional impact on the deformation, or stretching, of nearby fluid elements. Generally, fluid deformation is quantified by the Cauchy-Green strain tensor, $\mathcal{C}(\mathbf{x}_0)$, defined as,

$$\mathcal{C}(\mathbf{x}_0, t_0, t - t_0) = \left[\nabla \mathcal{F}_{t_0}^t(\mathbf{x}_0) \right]^T \nabla \mathcal{F}_{t_0}^t(\mathbf{x}_0).$$

$\mathcal{C}(\mathbf{x}_0, t_0, t - t_0)$ is a symmetric, positive definite tensor describing the strain undergone by a fluid particle dependent on its initial condition.

To define the FTLE field using this criterion, we wish to consider how a slight time-dependent perturbation to an initial point will evolve as it is acted upon by the flow map. To do this we consider a small perturbation $\delta\mathbf{x}(t_0)$ about the spatial coordinate vector $\mathbf{x}(t_0)$, so that at time t , $\mathbf{y}(t) = \mathbf{x}(t) + \delta\mathbf{x}(t)$ is the new position, with $\mathbf{x}(t)$ and $\mathbf{y}(t)$ being solutions to the governing equation (2.2). Using a Taylor series expansion for $\mathcal{F}_{t_0}^t(\mathbf{y}_0) = \mathcal{F}_{t_0}^t(\mathbf{y}(t_0))$ about $\mathbf{x}_0 = \mathbf{x}(t_0)$, we can write the deviation between the perturbed and original trajectories as,

$$\begin{aligned}\delta\mathbf{x}(t) &= \mathbf{y}(t) - \mathbf{x}(t) \\ &= \mathcal{F}_{t_0}^t(\mathbf{y}_0) - \mathcal{F}_{t_0}^t(\mathbf{x}_0) \\ &= \mathcal{F}_{t_0}^t(\mathbf{x}_0) + \nabla\mathcal{F}_{t_0}^t(\mathbf{x}_0)\delta\mathbf{x}(t_0) + \mathcal{O}\left(\|\delta\mathbf{x}(t_0)\|_2\right)^2 - \mathcal{F}_{t_0}^t(\mathbf{x}_0) \\ &= \nabla\mathcal{F}_{t_0}^t(\mathbf{x}_0)\delta\mathbf{x}(t_0) + \mathcal{O}\left(\|\delta\mathbf{x}(t_0)\|_2\right)^2,\end{aligned}$$

where $\mathcal{O}\left(\|\delta\mathbf{x}(t_0)\|_2\right)^2$ is in Landau notation, and can be assumed to be negligible as $\delta\mathbf{x}(t_0)$ is an infinitesimal. The magnitude of the perturbation at time t is then given by the L²-norm (Shadden et al. 2005),

$$\|\delta\mathbf{x}(t)\|_2 = \|\nabla\mathcal{F}_{t_0}^t(\mathbf{x}_0)\delta\mathbf{x}(t_0)\|_2.$$

Recall that the definition of a matrix norm may be written as $\|A\| = \max_{|\mathbf{x}|=1} \|A\mathbf{x}\| = \sqrt{\lambda_{\max}(A^T A)}$. To find the maximal stretching we consider the maximum L²- norm of the time-dependent perturbation,

$$\begin{aligned}\max_{\|\delta\mathbf{x}(t_0)\|=1} \|\delta\mathbf{x}(t)\|_2 &= \max_{\|\delta\mathbf{x}(t_0)\|=1} \|\nabla\mathcal{F}_{t_0}^t(\mathbf{x}_0)\delta\mathbf{x}(t_0)\|_2 \\ &= \sqrt{\lambda_{\max}\left([\nabla\mathcal{F}_{t_0}^t(\mathbf{x}_0)]^T \nabla\mathcal{F}_{t_0}^t(\mathbf{x}_0)\right)} \\ &= \sqrt{\lambda_{\max}\left(\mathcal{C}(\mathbf{x}_0, t_0, t - t_0)\right)},\end{aligned}$$

where the second line uses the definition of the matrix norm and $\lambda_{\max}\left(\mathcal{C}(\mathbf{x}_0, t_0, t - t_0)\right)$ is the maximum eigenvalue of the Cauchy-Green strain tensor evaluated at \mathbf{x}_0 . The quantity $\max_{\|\delta\mathbf{x}(t_0)\|} \|\delta\mathbf{x}(t)\|_2$ can then be rewritten as a maximum over a ratio of norms,

$$\begin{aligned}\max_{\delta\mathbf{x}(t_0)} \frac{\|\delta\mathbf{x}(t)\|_2}{\|\delta\mathbf{x}(t_0)\|_2} &= \sqrt{\lambda_{\max}\left(\mathcal{C}(\mathbf{x}_0, t_0, t - t_0)\right)} \\ &= e^{\text{FTLE}(\mathbf{x}, t_0, t) |t - t_0|},\end{aligned}$$

where $\delta\bar{\mathbf{x}}(t_0)$ is aligned with the eigenvector corresponding to $\lambda_{max}(\mathcal{C}(\mathbf{x}_0, t_0, t - t_0))$ and

$$\text{FTLE}(\mathbf{x}, t_0, t) = \frac{1}{|t - t_0|} \ln \left(\sqrt{\lambda_{max}(\mathcal{C}(\mathbf{x}_0, t_0, t - t_0))} \right), \quad (2.3)$$

is the FTLE field for the time interval $[t_0, t]$.

The FTLE gives the maximal stretching exponent for particles with initial position $\mathbf{x}(t_0) = \mathbf{x}_0$ over the time interval $[t_0, t]$.

How the FTLE gives a measure of exponential stretching can be illustrated simplistically with the steady linear system,

$$\dot{\mathbf{x}} = A\mathbf{x}. \quad (2.4)$$

The system (2.4) would give the flow map as,

$$\mathcal{F}_{t_0}^t(\mathbf{x}_0) = e^{At}\mathbf{x}_0,$$

such that the perturbation evolution is,

$$\delta\mathbf{x}(t) = e^{AT}\delta\mathbf{x}(t_0).$$

The FTLE would then be,

$$\text{FTLE}(\mathbf{x}, t_0, t) = \|A\|,$$

which measures the stretching induced by the largest eigenvalue of A .

If the FTLE is plotted as a field over a grid of initial positions then regions of conspicuously high FTLE values, often referred to as ‘ridges’, indicate regions of exceptional stretching. Ridges generally indicate barriers between distinct regions on the flow, allowing the partitioning of the flow into areas of differing structural characteristics. Ridges were originally used as a diagnostic for structures in oceanographic and atmospheric studies (Pierrehumbert & Yang 1993, von Hardenberg et al. 2000), although much research since has solidified the connection between FTLE ridges and extant coherent structures.

2.4 FTLE ridges and LCS

A crucial step in connecting a more rigorous LCS theory to the FTLE diagnostic was the recognition that FTLE ridges closely correspond to hyperbolic (saddle-like) points and their associated stable and unstable manifolds. As discussed in Section 2.1, in steady flows fluid particles asymptotically approach these manifolds in forward or backward time,

creating distinct barriers between regions of the flow. Similar behaviour occurs in idealised unsteady flows, although in this case the hyperbolic point becomes a hyperbolic trajectory moving through time, again with associated stable and unstable manifolds (Balasuriya 2016). The close connection between FTLE ridges and hyperbolic trajectories was first suggested by studies illustrating the close correspondence between FTLE contours and unstable manifolds of infinite time dynamical systems (Doerner et al. 1999). In meteorological applications, FTLE plots over initial conditions had also previously been investigated as indicators of barriers between mixing regions in the exosphere (Pierrehumbert & Yang 1993) and between transient eddy structures in baroclinic dynamics (von Hardenberg et al. 2000).

Following these early studies, Haller (2001) set out analytic conditions necessary for ridges and repelling and attracting material surfaces to coincide. Shadden et al. (2005) followed with a more formal definition of the method which made the FTLEs' relationship to hyperbolic LCS explicit. Their study shows that FTLEs represent nearly invariant manifolds even in systems with arbitrary time dependence. Later developments have further solidified the connection between repelling LCS and FTLE fields. Several studies (Haller 2011, Haller & Beron-Vera 2012, Karrasch 2012) use variational theory to give criteria which must be satisfied for FTLE ridges and hyperbolic LCS to coincide. More recently Balasuriya et al. (2016) demonstrate that particles traveling within hyperbolic neighbourhoods, regions of hyperbolic trajectories close to hyperbolic points, experience exponential stretching. They show that a greater time spent within such regions directly correlates to larger FTLE values, justifying the use of ridges as indicators of hyperbolic structures.

The FTLE is now one of the most popular techniques for identifying hyperbolic LCS in a variety of applications. For example, FTLE ridges have been used frequently to assist in the tracking and control of pollutant releases (Lekien et al. 2005, 2007). Indeed, a study by Mezic et al. (2010) uses the method to track the spread of oil in the infamous Gulf Oil Spill. Other interesting applications include a study which demonstrates how top marine predators often track oceanic FTLE structures to find food (Kew et al. 2009), studies using the technique to analyse sea-ice dynamics (Szanyi et al. 2016) and studies investigating how FTLE ridges impact the transport of airborne microorganisms (Bozorgmagham et al. 2013). A related technique called the Finite Size Lyapunov Exponent (FSLE) is also very common in the oceanographic and atmospheric literature. Rather than finding the maximal separation of fluid particles over a finite time, the FSLE method effectively finds the minimal time required for fluid particles in a flow to reach a particular magnitude of separation (Cencini & Vulpiani 2013). FSLE ridges were shown by Karrasch & Haller (2013) to coincide with FTLE ridges given certain conditions, and hence with hyperbolic LCS given Haller's further requirements (2001). The diversity of

disciplines which now commonly use the FTLE and FSLE are a testament to the relative ease of numerical implementation of these methods and their ability to capture a range of significant hyperbolic LCS from Lagrangian data.

2.5 Issues with FTLEs

Although one of the most widely used numerical techniques for the extraction of LCS, the use of FTLE ridges as indicators of hyperbolic LCS is only a heuristic unless certain criteria are satisfied (Haller 2011, Haller & Beron-Vera 2012, Karrasch 2012). Studies by Haller (2011) and Banicki and Wiggins (2010) detail various simple, analytic cases where the FTLE may give misleading structural information, such as those in which ridges are regions of high shear but do not correspond to attracting or repelling manifolds. Balasuriya et al. (2016) also address this issue by showing the FTLE diagnostic is only reliable when the hyperbolic neighbourhoods associated with extant manifolds have a specific geometrical structure. The criteria required to avoid such misdiagnosis, however, may be difficult to ascertain in complex or non-analytic flows. Shadden et. al. (2005, 2007) attempt to resolve issues of misdiagnosis by defining an LCS as a second derivative FTLE ridge for a fixed integration time, although this idea is problematic as second derivative FTLE ridges do not exist in all flows (Norgard & Bremer 2012).

Issues of noise sensitivity may also significantly impact the method's accuracy. Olcay, Pottebaum and Krueger (2010) found that the error in FTLE values when Gaussian noise is added increases greatly in highly transient portions of the flow with high spatial gradients. This local sensitivity to noise of the FTLE has been tackled by alterations to the original FTLE algorithm, such as through adaptive meshes which track not just the flow map and its gradient, but also the Hessian, to provide higher order corrections to the flow map (Mirona et al. 2015). Although this higher order formulation is more computationally expensive, further alterations to the algorithm have increased its efficiency while maintaining gains in accuracy (Fortina et al. 2015).

Other significant issues with the FTLE include difficulties in dealing with limited data at domain boundaries. Because trajectory integration must be stopped at the domain boundary, the edges of the domain act as attractors and are often misidentified as ridges when they do not correspond to LCS. This was highlighted as an issue by Tang et. al. (2010), who conceive a finite-domain version of the FTLE which they apply successfully in a study of velocity data from aircraft landing at the Hong Kong International Airport.

To amend various issues with the FTLE in practice, Kuhn et. al. (2012) devise various benchmarks for evaluating the best FTLE algorithm to use for a particular application.

Additionally, Balasuriya (2020b) derives theoretical error estimates for the FTLE based on parameters such as spatial resolution, allowing a characterisation of the degree of uncertainty in identified ridges. Although the FTLE has various issues, due to its ease of implementation and relevance to a range of applications it has continued to be a popular choice for LCS extraction. More accurate methods for LCS extraction are certainly available, although often not as straight-forwardly implementable, computationally efficient or easily interpreted (Peacock & Dabiri 2010).

2.6 Numerical calculation of FTLE fields

The FTLE may be numerically computed using a gridding approach similar to that used in finite-difference schemes. For simplicity, we consider a single initial position in two dimensions, labeled (x_0, y_0) . A grid of perturbations in the x and y directions from this initial point is then generated, as shown in Figure 2.3: $\{(x_0 + \delta x, y_0), (x_0 - \delta x, y_0), (x_0, y_0 + \delta y), (x_0, y_0 - \delta y)\}$. This grid of perturbed points can be evolved forward in time to time t using a numerical advection scheme, so that the final points are $\{\mathcal{F}_{t_0}^t(x_0 + \delta x, y_0), \mathcal{F}_{t_0}^t(x_0 - \delta x, y_0), \mathcal{F}_{t_0}^t(x_0, y_0 + \delta y), \mathcal{F}_{t_0}^t(x_0, y_0 - \delta y)\}$. The generated grids of initial and final points can then be used to approximate the gradient of the flow map as

$$\nabla \mathcal{F}_{t_0}^t \approx \begin{bmatrix} \frac{x(t; t_0, (x_0 + \delta x, y_0)) - x(t; t_0, (x_0 - \delta x, y_0))}{|2\delta x|} & \frac{x(t; t_0, (x_0, y_0 + \delta y)) - x(t; t_0, (x_0, y_0 - \delta y))}{|2\delta y|} \\ \frac{y(t; t_0, (x_0 + \delta x, y_0)) - y(t; t_0, (x_0 - \delta x, y_0))}{|2\delta x|} & \frac{y(t; t_0, (x_0, y_0 + \delta y)) - y(t; t_0, (x_0, y_0 - \delta y))}{|2\delta y|} \end{bmatrix}, \quad (2.5)$$

where $x(t; t_0, \mathbf{x}_0) := [\mathcal{F}_{t_0}^t(\mathbf{x}_0)]_1$ is the first component of the flow map and $y(t; t_0, \mathbf{x}_0) := [\mathcal{F}_{t_0}^t(\mathbf{x}_0)]_2$ is the second component, for initial condition $\mathbf{x}_0 = (x_0, y_0)^T$ (Haller 2002, 2015). Using this approximation, the Cauchy-Green strain tensor can be approximated by simply finding $[\nabla \mathcal{F}_{t_0}^t]^T \nabla \mathcal{F}_{t_0}^t$ for the approximate $\nabla \mathcal{F}_{t_0}^t$ in (2.5), giving,

$$\begin{aligned} [\mathcal{C}((x_0, y_0); t - t_0)]_{(1,1)} &\approx \left([\nabla \mathcal{F}_{t_0}^t]_{(1,1)} \right)^2 + \left([\nabla \mathcal{F}_{t_0}^t]_{(1,2)} \right)^2. \\ [\mathcal{C}((x_0, y_0); t - t_0)]_{(1,2)} &\approx [\nabla \mathcal{F}_{t_0}^t]_{(1,1)} [\nabla \mathcal{F}_{t_0}^t]_{(1,2)} + [\nabla \mathcal{F}_{t_0}^t]_{(2,1)} [\nabla \mathcal{F}_{t_0}^t]_{(2,2)}. \\ [\mathcal{C}((x_0, y_0); t - t_0)]_{(2,1)} &\approx [\nabla \mathcal{F}_{t_0}^t]_{(1,1)} [\nabla \mathcal{F}_{t_0}^t]_{(1,2)} + [\nabla \mathcal{F}_{t_0}^t]_{(2,1)} [\nabla \mathcal{F}_{t_0}^t]_{(2,2)}. \\ [\mathcal{C}((x_0, y_0); t - t_0)]_{(2,2)} &\approx \left([\nabla \mathcal{F}_{t_0}^t]_{(2,1)} \right)^2 + \left([\nabla \mathcal{F}_{t_0}^t]_{(2,2)} \right)^2. \end{aligned}$$

The maximum eigenvalue of the approximate Cauchy-Green strain tensor is then found numerically and used to calculate the FTLE as in the formula (2.3). This is done for




Figure 2.3: Schematic of numerically approximating the flow map gradient via gridding. A grid is drawn around each fluid particle's initial position, using deformations in the x direction ($\pm\delta x$) and y direction ($\pm\delta y$). This grid is propagated forward from t_0 to t using the flow map. The original and final grid positions, and the magnitude of δx and δy , are used to approximate the flow map gradient over $[t_0, t]$.

every initial condition in the set of initial points, such that plotting the FTLE at each corresponding initial point gives an FTLE field over $U(t_0)$.

2.7 The FTLE for Rossby Wave Flow

As Lorenz demonstrated with his analysis of the Lorenz attractor, toy mathematical models are very useful in demonstrating structural properties of geophysical flows while still being reasonably straightforward to represent mathematically. In the case of unsteady systems, one such model is that for two-dimensional Rossby wave flow, representing an important type of geophysical flow predominantly occurring in the atmosphere. This simplistic model is convenient as it features significant stretching and rotational structures and may be easily interchanged from steady to unsteady flow via a tunable parameter.

Geophysical Rossby wave flow arises from large-scale horizontal mixing and is central to many of the processes which strongly impact the climate, oceans and atmosphere of Earth. The flow arises from Earth's rotation and curvature, and is often associated with horizontally extensive, meandering streams which flow around mesoscale eddies (Balasuriya 2016). The study of atmospheric Rossby wave dynamics is of particular importance

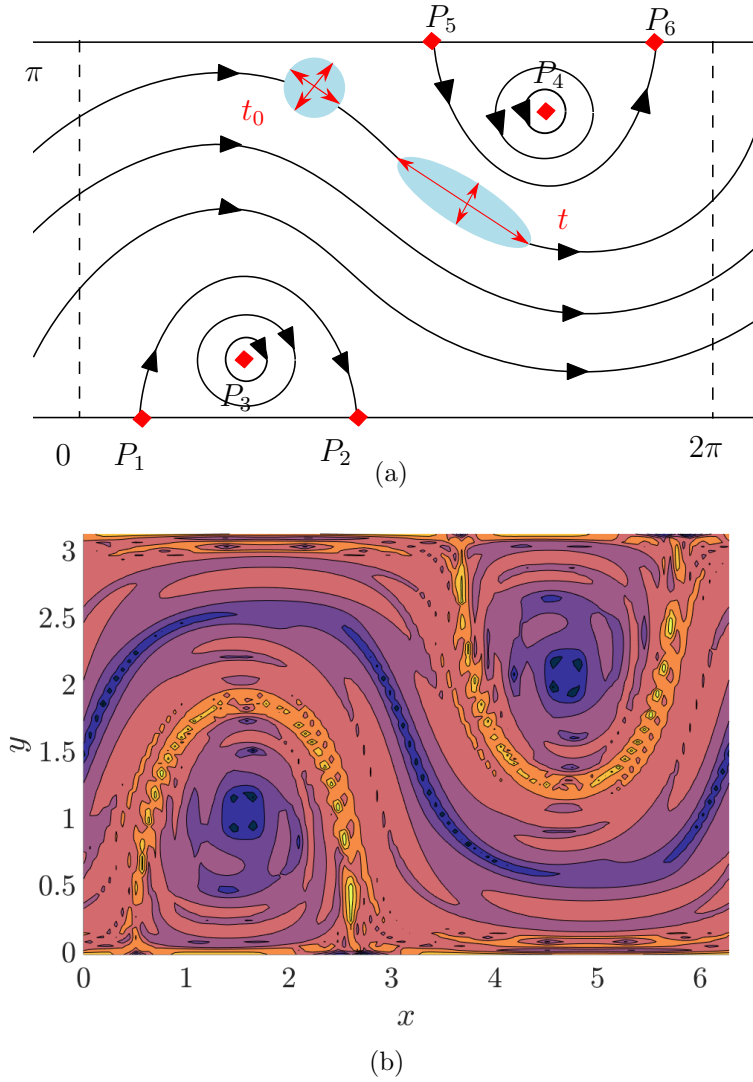


Figure 2.4: (a) shows the phase space diagram for Ψ in (2.8) with $\epsilon = 0$ in Equation 2.7. Saddle-type fixed points exist at P_1 , P_2 , P_5 and P_6 , tending to cause fluid particle stretching nearby. Centre fixed points are found at P_3 and P_4 , tending to cause fluid particle rotation nearby (Adapted from Pierrehumbert (1991a)). (b) shows an FTLE field for Ψ in (2.8) with $\epsilon = 0$ for $[t_0, t] = [0, 24]$. It can be seen that FTLE ‘ridges’ (yellow/orange) draw out boundaries between the stream region, where stretching dominates, and gyre regions, where rotation dominates.

in studies of climate as these are often responsible for planetary-scale energy transport (Wang et al. 2013). Disturbances in Rossby wave amplitude, for example, have been linked to sudden warming of the polar cap and break up of the circumpolar vortex (McIntyre & Palmer 1983). Rossby wave patterns have also been shown to have dominant statistical




Figure 2.5: FTLE fields for Ψ in (2.8) with $\epsilon = 0$ based on flow times (a) $[0, 3]$, (b) $[0, 6]$, (c) $[0, 12]$ and (d) for $[0, 24]$. From this sequential extension of the calculation time it can be seen how the FTLE ridge (orange/yellow) slowly extends to define the transport barrier between gyres and inner stream.

association with ground temperature changes in the Southern Hemisphere (Wang et al. 2013) and have been linked to several heat waves in the Northern Hemisphere (Petoukhov et al. 2013). More broadly, similar large-scale eddy structures are associated with the turbulent atmospheric dynamics of other rapidly rotating planets, such as Jupiter’s Great Red Spot and Neptune’s Great Dark Spot (Pierrehumbert 1991*b*).

Simple mathematical models of Rossby wave flow are useful for illustrating LCS extraction techniques as they demonstrate spatially traversing coherent structures while still possessing analytic flow velocity functions. The most simplistic mathematical model for Rossby wave type flow is the single wave model, which provides a good approximation to atmospheric meandering jets. Although quite simple, this form chooses waves which are dynamically consistent in that the relationship between the parameters obeys quasi-geostrophic dynamics, in which the Earth’s Coriolis force and horizontal pressure gradient forces are almost in balance. This dynamic consistency gives the model relevance for many geophysical applications (Pierrehumbert 1991*b*). The single wave model has the stream

function,

$$\Psi_0(x, y, t) = A \sin[k_0(x - c_0t)] \sin[l_0y], \quad (2.6)$$

where (x, y, t) are the local eastward and northward coordinates and time, and k_0 and l_0 are the wavenumbers for these directions respectively. The quantity $c_0 = -\beta/(k_0^2 + l_0^2)$ is the wavespeed, where β , the Rossby parameter, describes the variation in the Coriolis force due to Earth's spherical shape (Samelson & Wiggins 2006). This stream function is that of a single Rossby wave with wavelength $2\pi/k_0$ which propagates eastward at speed c_0 , with any northward meanders being constrained by the value of l_0 .

A slightly more realistic extension of the model in equation (2.6) is that in which a small time-periodic perturbation is added to the stream function, so that

$$\Psi(x, y, t) = \Psi_0 + \epsilon\Psi_1 = A \sin(k_1(x - c_1t)) \sin(l_1y) + \epsilon \sin(k_2(x - c_2t)) \sin(l_2y), \quad (2.7)$$

for small ϵ and where $c_i = -\beta/(k_i^2 + l_i^2)$ for $i = 1, 2$ is the wavespeed for the dominant and perturbing waves respectively. When this small wave perturbation is added, the quasi-geostrophic dynamics of the single-wave model Ψ_0 are obeyed to leading-order in the small parameter ϵ . For the stream function Ψ , the velocity field is

$$\begin{aligned} u &= -\frac{\partial\Psi_1}{\partial y} = -Al_1 \sin(k_1(x - c_1t)) \cos(l_1y) - \epsilon l_2 \sin(k_2(x - c_2t)) \cos(l_2y) \\ v &= \frac{\partial\Psi_1}{\partial x} = Ak_1 \cos(k_1(x - c_1t)) \sin(l_1y) + \epsilon k_2 \cos(k_2(x - c_2t)) \sin(l_2y). \end{aligned} \quad (2.8)$$

The quantity $(u(x, y, t), v(x, y, t))$ then represents x - and y - components of the flow velocity at spatial coordinates x and y and time t .

The model (2.7) is also commonly considered in a frame which moves with the first wave component of the flow, referred to as the co-moving frame (Samelson & Wiggins 2006). In this case we consider the transformation,

$$x^* = x - c_1t, \quad y^* = y,$$

so that the velocity field in the co-moving frame is,

$$\begin{aligned} u^* &= \frac{dx^*}{dt} = \frac{dx}{dt}(x^*, y^*) - c_1 \\ &= -Al_1 \sin(k_1x^*) \cos(l_1y^*) - \epsilon l_2 \sin(k_2(x^* - (c_2 + c_1)t)) \cos(l_2y^*) - c_1. \\ v^* &= \frac{dy^*}{dt} = \frac{dy}{dt}(x^*, y^*) \\ &= Ak_1 \cos(k_1x^*) \sin(l_1y^*) + \epsilon k_2 \cos(k_2(x^* - (c_2 + c_1)t)) \sin(l_2y^*). \end{aligned}$$

The streamlines of the stream function Ψ in the co-moving frame are shown in Figure 2.4a. The co-moving frame is advantageous as in this frame the streamlines of Ψ_0 coincide with those of Ψ when $\epsilon = 0$ and completely describe the Lagrangian motion of the flow for all times. I.e. for $\epsilon = 0$ the flow Ψ is steady, characterised by impermeable boundaries separating rotationally-dominated upper and lower gyres and a shear-dominated central stream. A steady phase space analysis of Figure 2.4a indicates the two rotational gyre structures in the presence of two centre fixed points, labeled P_3 and P_4 , and the central shear stream structure in four saddle-point type fixed points, labeled P_1 , P_2 , P_5 and P_6 (Pierrehumbert 1991a).

Comparing the FTLE field for Ψ with $\epsilon = 0$ over $[t_0, t] = [0, 12]$ in Figure 2.4b to the flow's steady phase space in Figure 2.4a, it can be appreciated how the FTLE field similarly designates distinct flow regions within which rotational or shear-type flow dominates. The FTLE ridges in yellow/orange define the transport boundaries between the outer gyres and inner stream, designating distinctly rotational and shear flow regions. Being an LCS method, however, the structural information given by the FTLE, even in the steady case, is dependent on the time interval over which the FTLE is calculated. Figures 2.5 (a), (b), (c) and (d) show the FTLE field for $\epsilon = 0$ with calculation times $[0, 3]$, $[0, 6]$, $[0, 12]$ and $[0, 24]$ respectively. From these plots it can be appreciated how the FTLE field paints out the evolving barriers in the flow as the calculation time is extended and fluid parcels become increasingly stretched by the flow.

In the unsteady case with $\epsilon > 0$, however, the rich structural information afforded by the FTLE field compared to simply plotting phase space trajectories can be appreciated. A comparison of Lagrangian trajectories for 36 fluid particles in the steady and unsteady case for Ψ is shown in Figure 2.6. As can be seen, in the steady co-moving frame case of Figure 2.6b, the particle trajectories perfectly match the structure of both the phase space diagram in 2.4a and steady FTLE field in 2.4b. In the the unsteady cases of 2.6c and 2.6d, however, the trajectories no longer give a clear picture of underlying flow structure. Here the trajectories of particles in the stream and gyre overlap and are difficult to separate. The barriers between the rotational gyres and central shear stream begins to break down, such that fluid initially in a gyre may now be pulled into the stream or vice versa. This ‘porosity’ of the gyre-stream barrier to the exchange of fluid scales with ϵ , and is best illustrated by considering FTLE fields (Samelson & Wiggins 2006).

FTLE fields for the unsteady flow with varied ϵ are shown in Figures 2.7a to 2.7d. As ϵ is increased from 0 in Figure 2.7a to 0.5 in Figure 2.7d the shape of the gyre boundaries becomes increasingly warped until they are nearly indiscernible from the inner stream. For $\epsilon = 0.1$ in Figure 2.7b, for example, we would expect a limited exchange of fluid between gyres and inner stream, as the ridges present in the $\epsilon = 0$ case are only slightly

warped and the overall flow structure retains its integrity. Comparatively, for $\epsilon = 0.3$ in Figure 2.7c the ridges present in 2.7a are significantly warped, such that the gyre boundaries are clearly porous to the inner stream. In Figure 2.7d this warping is increasingly extreme, such that parts of the FTLE field become a noisy expanse with small isolated structures scattered throughout, suggesting a highly chaotic regime. The parameter ϵ in Rossby wave flow is indeed closely linked to the degree of chaoticity of the flow, as it tunes the magnitude of the time-dependent perturbation to the steady flow. This chaoticity is in turn strongly connected to the magnitude of the FTLE field.

2.8 The parameter ϵ and the chaoticity of Rossby wave flow

The FTLE field is strongly related to a flow's degree of chaoticity as it measures the amount of stretching fluid parcels undergo as they are advected by the flow. As discussed in Chapter 1, a classic characteristic of chaos is exponential divergence of initial conditions over infinite time, often quantified by the Lyapunov exponent. The Lyapunov exponent is effectively the maximal rate of divergence between two points which are initially infinitesimally close in the limit as time approaches infinity. Thus, although the FTLE is calculated over finite times, there is a significant link between the FTLE and a flow's chaoticity when the calculation time is sufficiently large for FTLE ridges to form. For Rossby wave flow, the flow exhibits chaotic behaviour when $\epsilon > 0$, but the FTLE may be used to gain a more quantitative insight into how the degree of chaos in the flow scales with the magnitude of ϵ and the computation time, $[t_0, t]$.

As the FTLE has units of 1/time, the quantities $1/FTLE_{med}$ and $1/FTLE_{av}$, where $FTLE_{av}$ is the average FTLE over the flow domain and $FTLE_{med}$ is the median, may be considered a rough measure of the time scale over which a system becomes chaotic. Calculating $1/FTLE_{med}$ or $1/FTLE_{av}$ for $t_0 = 0$ and different integration times $T_f = t - t_0$ gives a measure of the timescale of chaotic behaviour in the system. The quantity $1/FTLE_{av}$ gives some measure of the average timescale of chaoticity for the set of initial conditions considered for the flow, suggesting how the magnitude of chaotic timescales is distributed over the set of initial conditions. The quantity $1/FTLE_{med}$ gives a slightly better estimate of the timescale for chaoticity as it corresponds to the most frequently occurring FTLE value, suggesting the most frequently occurring chaotic timescale over the initial condition set. The average will be impacted more by outliers than the median, and in highly chaotic situations the FTLE calculations contain more uncertainty, so using the median is a more robust approach. As the ϵ parameter scales the time-dependent perturbation to the flow, it is expected that it is closely linked to the degree of chaos in the system.

Motivated by these considerations, and that this relationship has not been investigated in this way previously, the quantities $1/FTLE_{med}$ and $1/FTLE_{av}$ are plotted for different ϵ and $T_f = 1, 2, 3, 6, 24, 36$ and 48 in Figure 2.8.

T_f	ϵ_{min}
12	0.6
24	0.55
36	0.47
48	0.45

Table 2.1: Table of the parameter ϵ_{min} giving the $1/FTLE_{av}$ minimum for each of the larger calculation times T_f .

As can be seen from Figure 2.8, for the smaller calculation times of $T_f = 1, 2, 3$ and 6 , both $1/FTLE_{med}$ (orange dashed) and $1/FTLE_{av}$ (blue fill) decay for increasing ϵ . This implies that for smaller calculation times the chaoticity only grows with ϵ . The plots indicate that the timescale for chaotic behavior is largest for ϵ near zero and smallest for $\epsilon > 1$. Comparatively, for the larger calculation times of $T_f = 12, 24, 36$ and 48 , minima in the $1/FTLE_{med}$ and $1/FTLE_{av}$ plots become apparent, suggesting that a particular $\epsilon = \epsilon_{min}$ gives a minimum timescale for chaoticity to occur. For ϵ greater than this ϵ_{min} , the timescale for chaoticity grows slightly before beginning to asymptote with increasing ϵ . Hence, unlike for smaller calculation times, there is no simple decay of the chaotic timescale with increasing ϵ and so the ϵ regime considered for a particular T_f becomes significant.

Table 2.1 details the ϵ_{min} giving an approximate minimum in $1/FTLE_{med}$ for each T_f . From Table 2.1 it can be seen that ϵ_{min} decreases with increasing T_f , suggesting that the ϵ giving the smallest timescale for chaotic behaviour decreases as T_f increases. This implies that large calculation times require a smaller magnitude of perturbation to the steady system for the highest possible chaoticity to occur than for lesser calculation times. Above a certain magnitude of perturbation however, the timescale of chaotic behaviour remains fairly constant, suggesting that for longer calculation times increasing the magnitude of perturbation to the steady system does not increase the chaoticity once a certain level of chaoticity is reached. Interestingly, for $T_f \leq 6$, $1/FTLE_{med}$ lies above $1/FTLE_{av}$, while for $T_f > 6$, $1/FTLE_{med}$ and $1/FTLE_{av}$ almost overlap. This suggests that for $T_f \leq 6$ the distribution of chaotic timescales over all the initial conditions in the domain is skewed towards higher than average chaotic timescales, while for $T_f > 6$ the distribution of chaotic timescales is fairly symmetric.

Intuitively, the chaotic timescale should also be linked to how well-defined FTLE ridges

are, as this is linked to the degree of exponential stretching of fluid particles over a finite time. The degree of exponential stretching of fluid particles will be impacted by the parameter ϵ , as this determines how quickly chaoticity occurs in the flow, as well as the time over which the model is run T_f , because, as shown above, this impacts how quickly chaoticity is reached. The relationship between the magnitude of ϵ , T_f and the formation of FTLE ridges is best illustrated by Figure 2.9, which shows a grid of FTLE fields with a selection of T_f along the y-axis and a selection of ϵ along the x-axis. As can be seen, when T_f is small, the FTLE ridges may be poorly defined for all ϵ , as there is insufficient time for particle stretching to occur, as is the case for $T_f = 0.5$. For some T_f , the chaoticity must be quite high for FTLE ridges to form, as in the case of $T_f = 2$, where ridges are less developed for $\epsilon = 0$ and $\epsilon = 0.3$, but better defined for higher ϵ . If T_f is sufficiently large, however, FTLE ridges form for most chaoticity regimes, as is the case of $T_f = 6, 12$ and 24 , for which ridges form for all ϵ . Interestingly, for $T_f > 6$, the chaotic timescale minimising ϵ identified in Table 2.1 seem to give a more noisy FTLE field, suggesting FTLE fields corresponding to such ϵ contain more uncertainty. For $\epsilon = 0.5$ for example, which is near the chaotic timescale minimising ϵ for $T_f = 12$ and $T_f = 24$, the FTLE fields seem to be undergoing a transition from a state in which the two gyres dominate the flow (for $\epsilon = 0.3$) to that in which multiple rotational eddies dominate the flow (for $\epsilon = 0.9$). This transition state is slightly noisy, containing elements of the two different structural scenarios, particularly in the $T_f = 24$ case. These images can only provide an intuition for regimes in which FTLE ridges will effectively form, but will assist in understanding results presented in Chapters 4 and 5.

Overall, these findings motivate the calculation times used in the investigations in Chapters 4 and 5, which use T_f values for FTLE calculations within 12 and 24. The range of ϵ values used in Chapters 4-5 is within $[0.1, 1.3]$, so that the shortest chaotic timescale (and hence highest chaoticity regimes) occurs in the middle range of these values. This is important as the aim of this thesis is to explore techniques which are robust in chaotic conditions. Using these timescales for this range of ϵ values thus provides some assurance that chaotic behaviour is occurring for the parameter regimes being used.

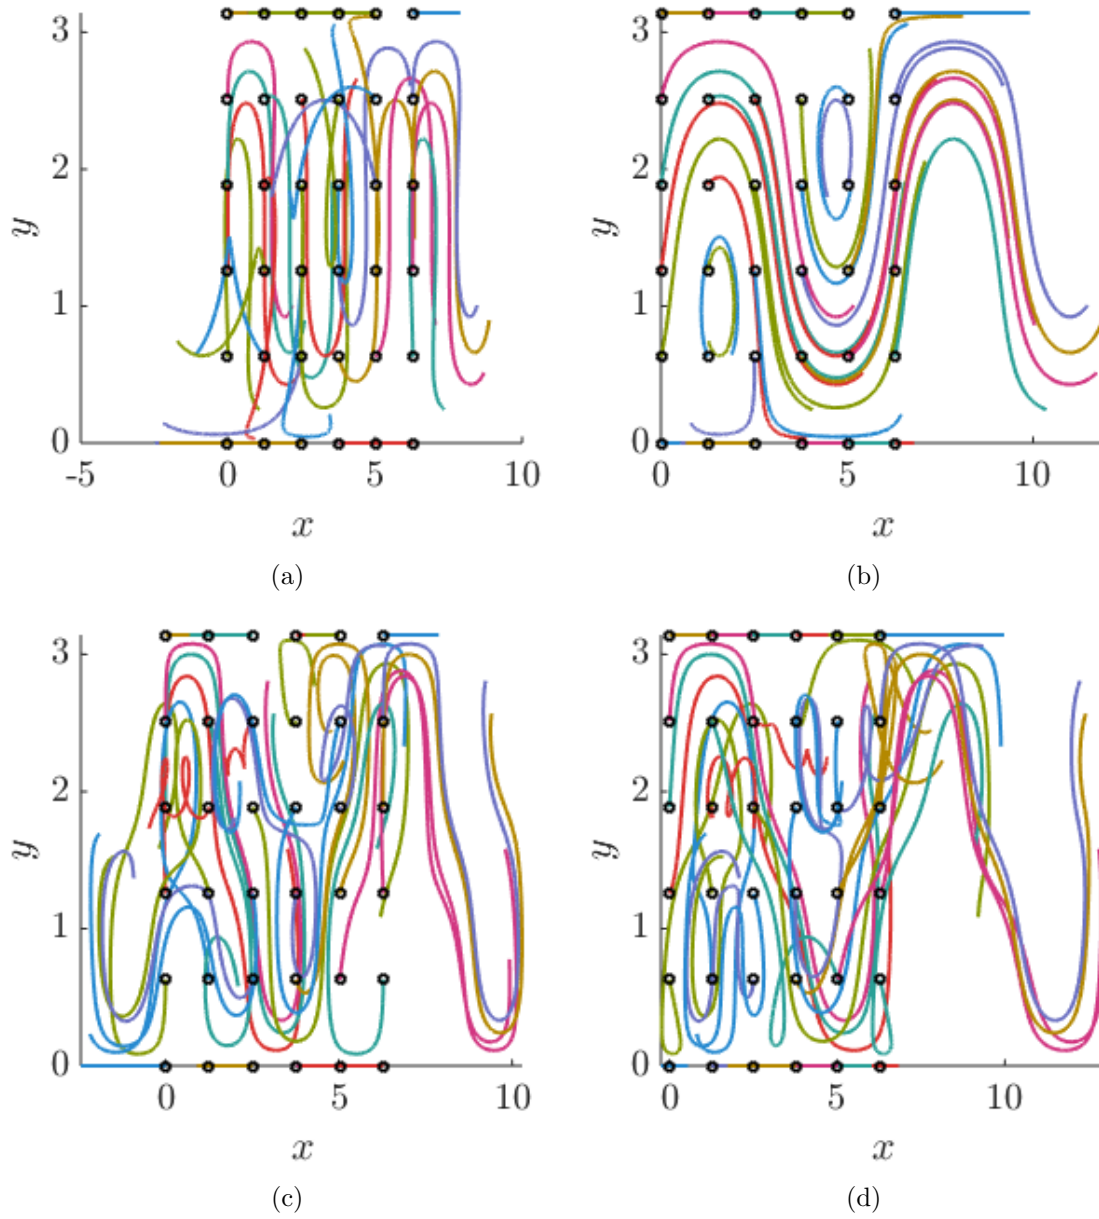


Figure 2.6: Lagrangian trajectories of 36 fluid particles initially placed on a uniform grid for the time $[t_0, t] = [0, 6]$. Different colours indicate trajectories of different initial positions (indicated with open black circles) on the grid. (a) for the stream function Ψ in the original frame with $\epsilon = 0$, (b) in co-moving frame with $\epsilon = 0$, (c) in the original frame with $\epsilon = 0.6$ and (d) in the co-moving frame with $\epsilon = 0.6$. It can be seen that for $\epsilon = 0$ in the co-moving frame the trajectories mirror the phase space diagram for the steady system shown in Figure 2.4a.




Figure 2.7: FTLE fields for $[t_0, t] = [0, 12]$ for (a) $\epsilon = 0$, (b) $\epsilon = 0.1$, (c) $\epsilon = 0.3$ and (d) $\epsilon = 0.5$. These small incremental increases in ϵ illustrate the sensitivity of the time perturbed Rossby wave model to changes in ϵ . Chaotic behaviour appears to occur more quickly for higher ϵ , a phenomenon discussed in detailed in Section 2.8.

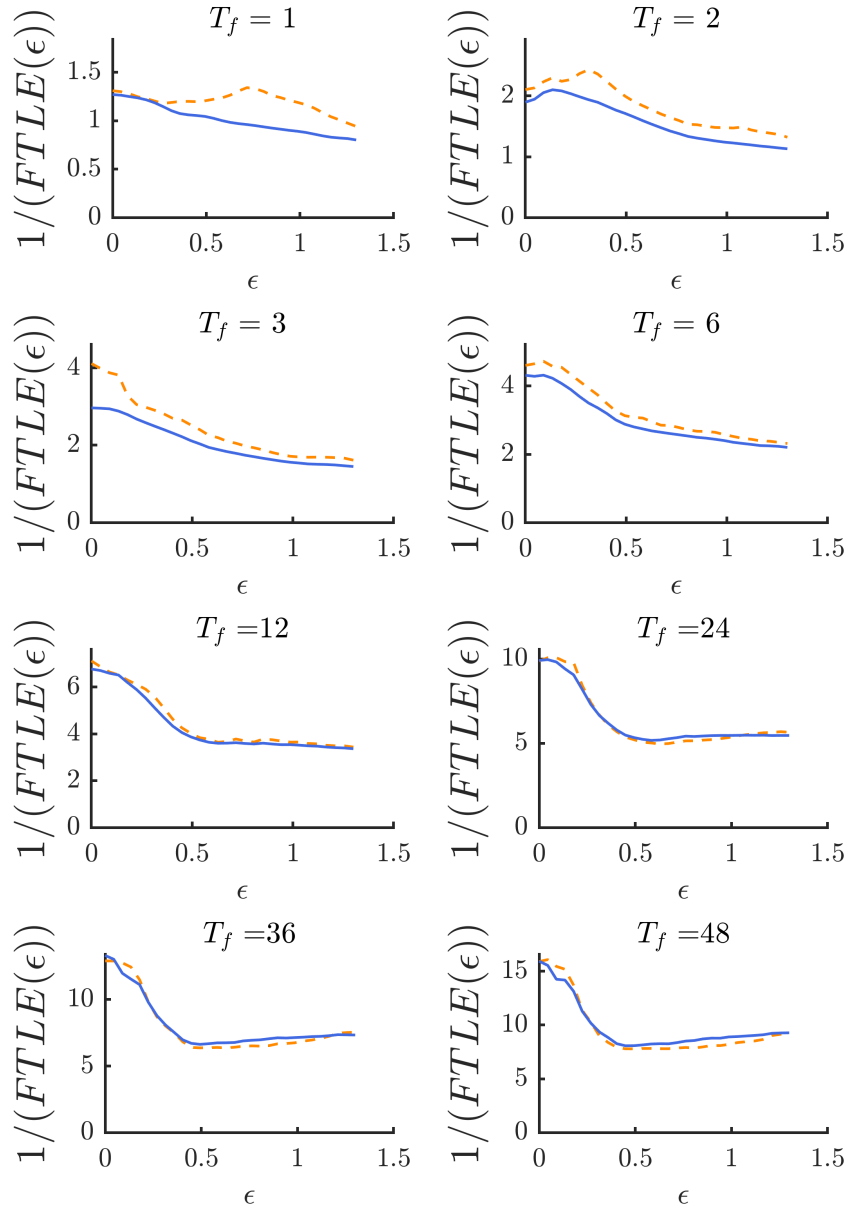


Figure 2.8: $1/FTLE_{med}$ (orange dashed) and $1/FTLE_{av}$ (blue fill) curves estimate the chaotic timescale for a range of ϵ and for different total calculation times T_f . As can be seen, for $T_f \leq 6$ the chaotic timescale decreases with increasing ϵ , while for $T_f > 6$ the chaotic timescale decreases until it is minimised by a particular ϵ_{min} .

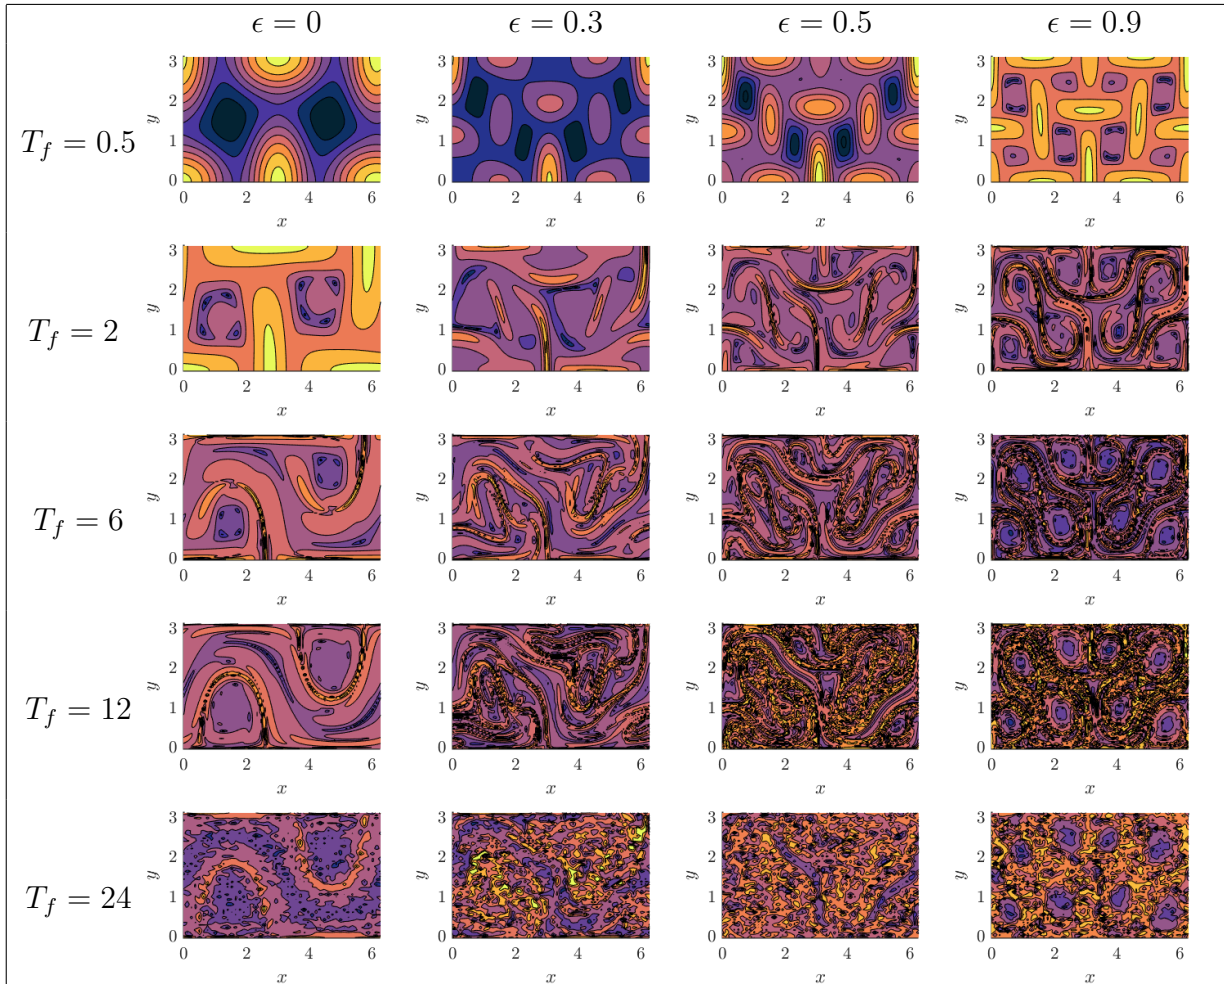


Figure 2.9: Table showing how the FTLE field varies with the calculation time T_f and the parameter ϵ .

Chapter 3

Data Assimilation

Data assimilation (DA) is a dominant methodology for making predictions about geophysical systems that seeks to optimally combine information from observational data and mathematical models. In particular, DA may be used to estimate geophysical model parameters, which is crucial for running models of systems such as the climate. Chapter 3 gives an overview of DA for geophysical systems and estimating geophysical model parameters. This builds a basis for discussing the DA algorithms which will be used to develop an FTLE-based data assimilation method for parameter estimation. In particular, two basic DA algorithms for parameter estimation, which will be used as benchmark comparisons for the new algorithm, are discussed in detail. Current research in LCS-based data assimilation methods is also reviewed to establish key areas for investigation for the algorithm developed in this thesis.

3.1 Data Assimilation for Geophysical Systems

Generally two forms of information are available to supplement a prediction about a geophysical system: observational data, such as temperature, wind speed and salinity, and a mathematical model, built from an understanding of the physics, chemistry and biology of a system. Both forms of information have their benefits and flaws. Observational data is a direct measurement of the system state, yet it contains measurement and instrumentation errors and, as mentioned in Chapter 1, is often distributed irregularly. Mathematical models can be used to fill these spatio-temporal gaps in data, yet, due to the complexity and high-dimensionality of geophysical systems these are highly simplified versions of the true system. Data Assimilation aims to combine the information captured from both these sources in a way which exploits their respective benefits yet minimises their downfalls (Lahoz & Schneider 2014) .

An example of the data assimilation process in a geophysical context is shown in Figure 3.1, which illustrates a DA scheme for making a prediction about global ozone spread. Ozone data, which possesses associated error due to the instrumentation used for data collection and other factors, is collected via satellite (shown at left in Figure 3.1). This collected data is often referred to as the *observation*. A mathematical model is also used to create a prediction for the global distribution of ozone (at right in Figure 3.1), which similarly possesses associated error due to model assumptions and other factors. This physical model is referred to as the *forecast*, and is generally of a higher resolution than the data. The data assimilation process combines the forecast and observation to produce what is called the *analysis* (at bottom in Figure 3.1), an estimate of the current state of the system. The analysis also possesses error, which may arise partially due to the errors in the observations and forecast, as well as due to the technique used to combine the two (Kalnay 2002). The key aim of data assimilation is to combine the observation and forecast such that the errors present in either are optimally balanced, minimising the error in the analysis. The specific process used to balance observation and forecast may be drawn from a range of mathematical disciplines, including optimisation, dynamical systems theory and Bayesian statistics.

Data assimilation for geophysical modeling predominantly arose due to researchers attempting to amend error issues in numerical weather prediction (NWP). The problem faced was well summarised by the meteorologist Charney, who in a 1951 review of NWP comments that “Owing to their statistically indeterminate behaviour, the turbulent properties of the atmosphere place an upper limit to the accuracy obtainable by dynamical methods of forecasting, beyond which we shall have to rely upon statistical methods.” (Charney 1951). By turbulent, indeterminate behaviour, Charney meant the innately chaotic properties of the atmosphere. Essentially, both meteorological data, analyzable using statistical methods, *and* model forecasts are needed to overcome the limits of predictive models alone.

NWP is essentially an initial value problem, as initial values are needed to run weather simulations into the future. An issue faced in the early history of geophysical data assimilation was how to use collected atmospheric data, which was not necessarily located at model grid points, to update states at model grid points. One might consider simply interpolating the observations to the grid points, however, often the data available are not sufficient if only interpolative methods are used. For example, NWP models may have degrees of freedom of the order 10^7 . Each grid point requires at least 4 initialisation values: 2 horizontal wind components, temperature and moisture, such that millions of variables must be initialised. However, for a time window of 3 hours, typically only around a million observations of the atmosphere are collected, far less than the number of degrees of freedom of the model (Kalnay 2002). Additionally, as mentioned in Chapter

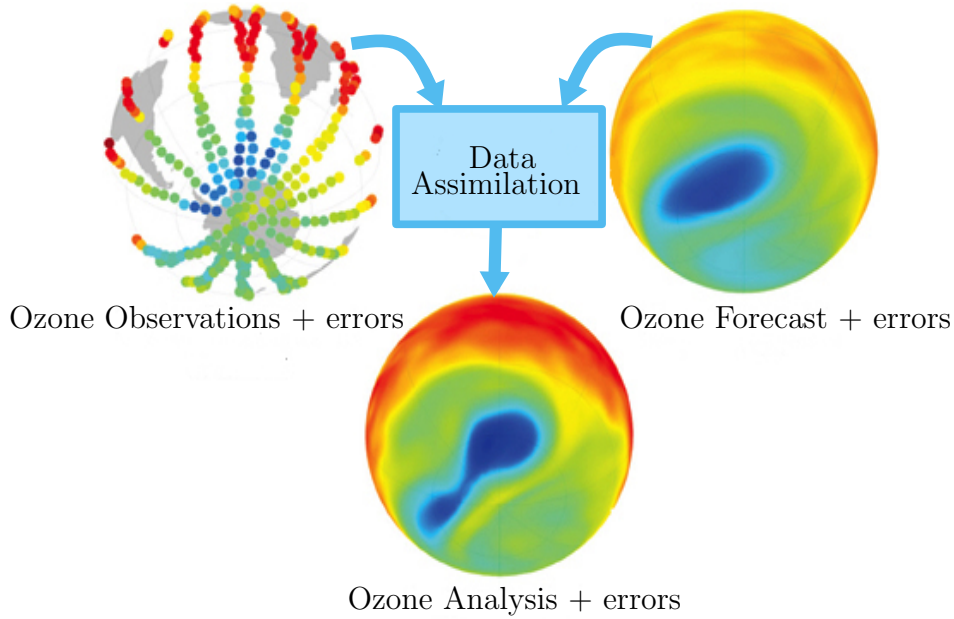


Figure 3.1: A schematic of how data assimilation combines information from a mathematical model with collected observational data. Ozone data (an *ozone observation*) is combined with a mathematical model for the geographical distribution of ozone (an *ozone forecast*). Both observations and forecast contain errors associated with data collection methods and model assumptions respectively. The forecast and observations are combined using a data assimilation scheme to form the *ozone analysis*, a prediction for global ozone levels at the time of the forecast and observation. This analysis also has associated error arising from a combination of measurement error, model assumptions and the data assimilation scheme itself. (Adapted from Figure 2 in Lahoz & Schneider (2014))

1, these data are often erratically distributed, due to factors such as satellite paths and cloud cover, as well as some parts of the Earth being relatively data-poor in comparison to America and Eurasia.

To amend this mismatch between high-dimensional models and lower dimensional, sparsely distributed data, DA takes a different approach to the interpolation problem. Instead of attempting to interpolate the observations $\mathbf{y}_n \in \mathbb{R}^N$ at time t_n to the model forecasts $\mathbf{x}_n \in \mathbb{R}^M$ (where $N \leq M$), the model forecasts are interpolated to the lower dimensional observation space. This conversion of \mathbf{x}_n from model variables to observed variables is generally described by the observation operator $h : \mathbb{R}^M \rightarrow \mathbb{R}^N$, which acts

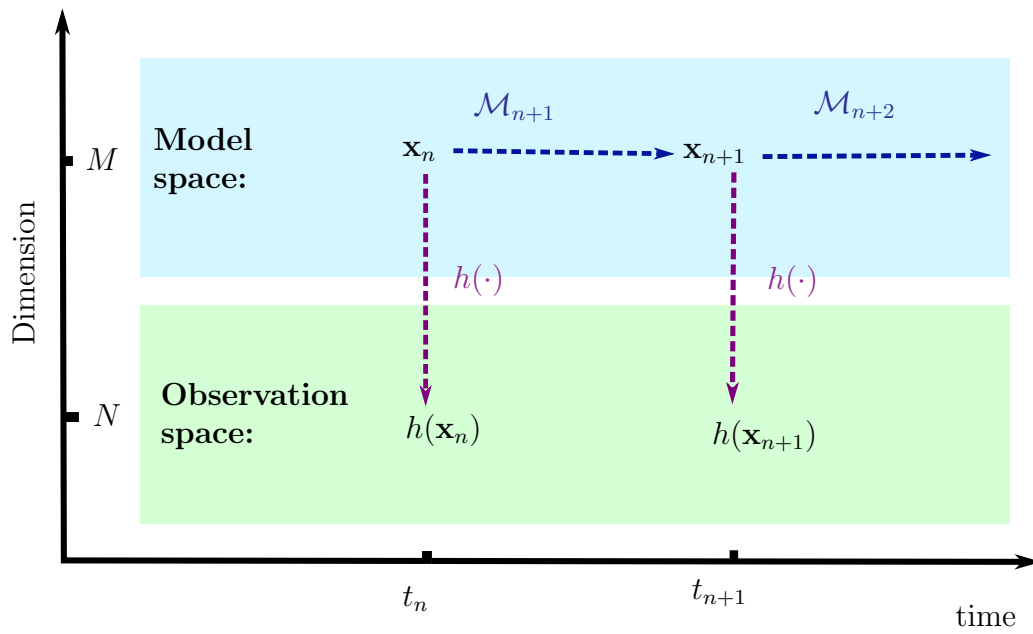


Figure 3.2: Diagram showing the operation of the observation operator $h(\cdot)$ and the model update \mathcal{M}_{n+1} in a generic DA scheme. The observation operator transfers the model state vector at time t_n , \mathbf{x}_n , from the model space, of dimension M , to observation space, of dimension $N \leq M$. Formulation of the analysis update to \mathbf{x}_n is then carried out in observation space. \mathbf{x}_{n+1} is then calculated by evolving the analysis update of the state forward in time using \mathcal{M}_{n+1} . This is repeated at each assimilation step.

upon \mathbf{x}_n , giving $h(\mathbf{x}_n)$, the model forecast in observation space (Law et al. 2015, Kalnay 2002). Computations to formulate the analysis are then carried out in the lower dimensional observation space, before transforming back to model space to calculate the model update. A schematic detailing this process for a generic DA scheme is shown in Figure 3.2. This approach to the problem amends many of the dimensional issues associated with interpolation, as well as reducing computational expense.

The defining quality of a particular data assimilation scheme is its update step at each time t_n ; how it combines observations \mathbf{y}_n and forecasts \mathbf{x}_n to produce the analysis. There are three general approaches to producing the analysis: the Kalman Gain, Variational and Bayesian approaches. Kalman Gain approaches, such as the Kalman Filter described in Section 3.2, formulate the update as a linear combination of \mathbf{x}_n and \mathbf{y}_n , where the weighting of terms in the linear combination is calculated via a minimisation of the difference in observation and forecast. Variational approaches, although not used in this thesis, balance information from the data and model via a cost function, which penalises disparities in the model and observations. Bayesian approaches, such as particle filters (described in Section 3.5) and Approximate Bayesian Computation (described in Section 3.6), differ from these first two approaches in that, rather than simply updating the system state in the analysis step, the probability distribution of the system state is updated (Law et al. 2015). This is useful as various summary statistics may then be calculated from the distribution, such as expected values, variances and prediction intervals.

3.2 DA Methods

As the models and data being dealt with in DA are generally time-dependent, the analysis is typically updated at regular intervals over time, called the *assimilation window*. At each time t_n the analysis from the previous time \mathbf{x}_{n-1}^a , where n denotes the assimilation window index, is updated to the current time using the model update equations (3.1). The update is denoted \mathbf{x}_n^p , with the p standing for prediction:

$$\mathbf{x}_n^p = \mathcal{M}_n(\mathbf{x}_{n-1}^a) + \beta_n, \quad (3.1)$$

where $\mathcal{M}_n : \mathbb{R}^M \rightarrow \mathbb{R}^M$ is the deterministic model from time t_{n-1} to t_n and β_n represents model error. It is generally assumed that the distribution of β_n is known. At each time an observation \mathbf{y}_n is also collected, which is related to the true state of the system \mathbf{x}_n^{true} via,

$$\mathbf{y}_n = h(\mathbf{x}_n^{true}) + \xi_n, \quad (3.2)$$

where ξ_n is random noise representing measurement errors and potentially representational errors if a truncated version of h is used (Leeuwen et al. 2019). One of the foundational

methods of data assimilation, which provides a useful illustration of how the analysis may be formulated from the forecast \mathbf{x}_n^p and observations \mathbf{y}_n , is the Kalman Filter, the most simple of the Kalman Gain methods (Evensen 2009). In this method Gaussian distributed error in the model and observations is assumed, so that,

$$\begin{aligned}\xi_n &\sim \mathcal{N}(\mathbf{0}, \mathbf{R}_n), \\ \beta_n &\sim \mathcal{N}(\mathbf{0}, \mathbf{Q}_n),\end{aligned}\tag{3.3}$$

for observational covariance matrix \mathbf{R}_n and model covariance matrix \mathbf{Q}_n at time t_n . A linear observation operator is also assumed, so that,

$$\mathbf{y}_n = H\mathbf{x}_n^{true} + \xi_n\tag{3.4}$$

where $H \in \mathbb{R}^{N \times M}$. At each time t_n , the new analysis is obtained via the Kalman update equations,

$$\mathbf{x}_n^a = \mathbf{x}_n^p + K_n [\mathbf{y}_n - H\mathbf{x}_n^p]\tag{3.5}$$

The Kalman gain matrix, $K_n \in \mathbb{R}^{M \times N}$ at time t_n , is a weight matrix derived from finding the solution to the least squares minimisation of the *innovation*, $\mathbf{y}_n - H\mathbf{x}_n^p = I_n$, which gives a measure of how much the observed data \mathbf{y}_n differs from the model's prediction \mathbf{x}_n^p (Law et al. 2015). The matrix K_n preferences model variables which are closest to the observational data points in terms of a Euclidean distance. The basic process of the Kalman Filter is shown in Figure 3.3, illustrating how the prediction, update and analysis steps move between model space and observation space to exploit the lower dimensionality of observation space .

A drawback of the regular Kalman filter, however, is that it requires a linear model and observation operator. Many geophysical systems are highly non-linear and hence DA for geophysical applications must be compatible with this characteristic non-linearity. Due to this, many extensions of the Kalman Filter have been derived to accommodate non-linearity, the most common of which are the extended Kalman Filter and the ensemble Kalman Filter (EnKF). The extended Kalman Filter is the most simplistic extension of the Kalman Filter in that it is basically the original Kalman Filter using a linearised observation operator, in Equation (3.4), and model update equation (3.1). The evolution of the error covariance in this approximation, however, is often unstable, rendering the method unsuitable for many applications. In addition to this issue, for a large model dimension the covariance may become too large to be realistically computable. Due to these problems, the Ensemble Kalman Filter is more commonly used for geophysical DA (Leeuwen 2015).

In the EnKF a Monte Carlo approach is taken to the regular Kalman Filter, with the analysis distribution at each time approximated via an ensemble of N state vectors.

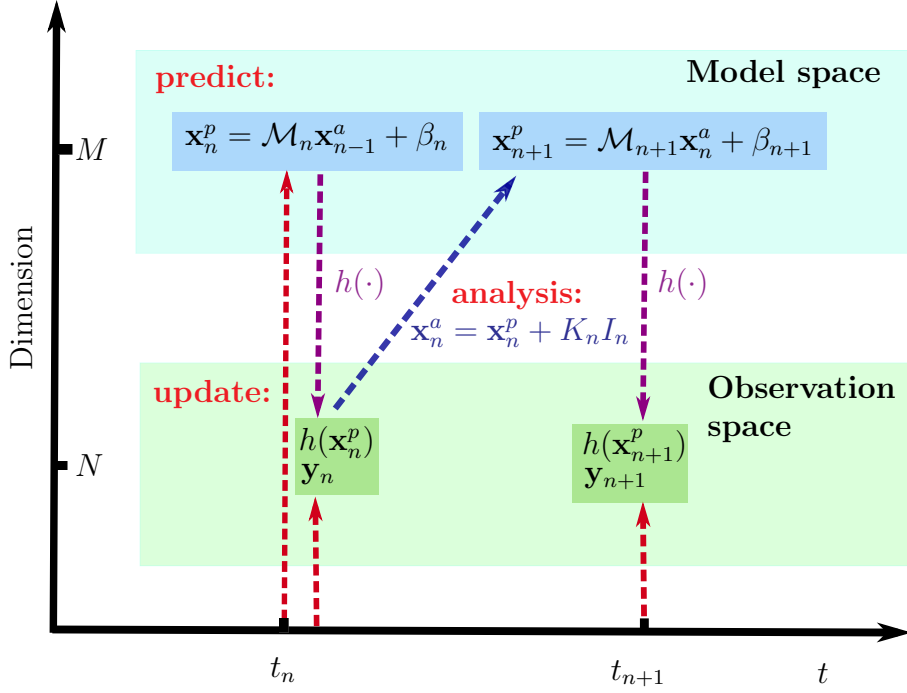


Figure 3.3: Schematic showing one assimilation step of the Kalman filter. A prediction \mathbf{x}_n^p at time t_n is produced by evolving the analysis from the previous time step forward in time using the model \mathcal{M}_n . The quantity \mathbf{x}_n^p is transferred to observation space using the observation operator $h(\cdot)$. The model in observation space $h(\mathbf{x}_n^p)$ is compared with \mathbf{y}_n , data collected at time t_n , via the innovation $I_n = \mathbf{y}_n - h(\mathbf{x}_n^p)$. The analysis \mathbf{x}_n^a is then formulated in model space by acting K_n , the Kalman matrix, on I_n , which weights the elements of I_n while also transferring I_n to model space. The process is repeated for each time step in the assimilation scheme.

Each ensemble member evolves according to the full non-linear model equations, thus ameliorating the instability issues associated with simply linearising the non-linear model (Vetra-Carvalho et al. 2018). The analysis covariance is then updated as the covariance of the ensemble, while several conditions assure that the mean analysis of the ensemble satisfies the Kalman update equation (3.5) (Evensen 2009). The EnKF is also a popular choice for parameter estimation in geophysical applications, where both the state *and* parameters of a system must be estimated (Katzfuss et al. 2016).

3.3 Bayesian Parameter Estimation

An important problem in many geophysical modeling applications is that of parameter estimation, often called the ‘inverse problem’. While at times it may be of interest to predict the future state of a system given particular initial inputs (the ‘forward problem’), often it is desirable to estimate the input parameters which would result in the observed data. This is a common scenario in the modeling of climate, oceans and atmosphere, as the models used often take in multiple parameters which must first be estimated from data (Villagran et al. 2008).

Due to chaotic behaviour, key parameters may significantly impact model predictions about the future state of a system, making it desirable to develop methods which estimate such parameters with considerable accuracy. This is particularly the case with climate models, for which even relatively small uncertainties in parameters may induce large uncertainties in climate projections (Villagran et al. 2008). In estimating such parameters it is desirable to also estimate their associated error so that this can be used to ascertain the impact of estimation error on model predictions. Thus, statistical methods are common in parameter estimation, as they allow various error statistics associated with an estimate to be straightforwardly calculated .

As parameters are not directly observable, parameter estimation can only be carried out in conjunction with state estimation, an approach referred to as *joint state-parameter estimation*. In the Bayesian statistical framework, the ultimate aim is to estimate the posterior distribution $p(\theta_n, \mathbf{x}_n | \mathbf{y}_n)$ for the parameter θ_n , model states \mathbf{x}_n and observations \mathbf{y}_n at time t_n . This is a probability density function (pdf) that gives the joint probability of a model state \mathbf{x}_n and parameter θ_n given that \mathbf{y}_n has been observed. The posterior concentrates probability density around the likely true state and parameter of the system at time t_n (Vetra-Carvalho et al. 2018, Law et al. 2015). Other statistical objects which form the basis of parameter estimation methods are the joint state-parameter prior $p(\theta_n, \mathbf{x}_n)$ and likelihood $p(\mathbf{y}_n | \theta_n, \mathbf{x}_n)$. The prior is a pdf selected based on beliefs about θ_n before any analysis is done, which may arise from consideration of the physical meaning of θ_n . The likelihood is a pdf which describes the probability of observing \mathbf{y}_n given the joint model-parameter state is known to be (θ_n, \mathbf{x}_n) . The likelihood and prior are key mathematical objects in parameter estimation because they are related to the target posterior $p(\theta_n, \mathbf{x}_n | \mathbf{y}_n)$ through Bayes’ Law:

$$p(\theta_n, \mathbf{x}_n | \mathbf{y}_n) = \frac{p(\mathbf{y}_n | \theta_n, \mathbf{x}_n)p(\theta_n, \mathbf{x}_n)}{p(\mathbf{y}_n)}, \quad (3.6)$$

where $p(\mathbf{y}_n)$ is a pdf representing the probability of observing \mathbf{y}_n at time t_n .

With the aim of formulating the posterior from Bayes' Law, many parameter estimation methods assume Markovian model states, such that each step in the assimilation scheme is conditioned only on the previous state for all variables and parameters (Kunsch 2013). In the Bayesian framework, this Markovian assumption means that the probability of a particular model state \mathbf{x}_l at time t_l , when conditioned on all previously occurring states, only depends on the time t_{l-1} . If $p(\mathbf{x}_{0:T}, \theta_{0:T})$ is the joint probability of a series of model-parameter states (\mathbf{x}_l, θ_l) from times $t_l = t_0$ to $t_l = t_T$, then the Markovian assumption requires that

$$p(\mathbf{x}_{0:T}, \theta_{0:T}) = p(\mathbf{x}_0, \theta_0) \prod_{l=1}^T p(\mathbf{x}_l, \theta_l | \mathbf{x}_{l-1}, \theta_{l-1}). \quad (3.7)$$

The observations \mathbf{y}_n are also generally assumed to be conditionally independent in time, so that

$$p(\mathbf{y}_{1:T} | \mathbf{x}_{0:T}, \theta_{0:T}) = \prod_{l=1}^T p(\mathbf{y}_l | \mathbf{x}_l, \theta_l), \quad (3.8)$$

where $\mathbf{y}_{1:T}$ is the set of observations collected over times t_1 to t_T . Using (3.7) and (3.8) within Bayes' Law (3.6), we can write

$$\begin{aligned} p(\mathbf{x}_T, \theta_T | \mathbf{y}_T) &= \frac{p(\mathbf{y}_{1:T} | \mathbf{x}_{0:T}, \theta_{0:T})}{p(\mathbf{y}_{1:T})} p(\mathbf{x}_{0:T}, \theta_{0:T}) \\ &\propto p(\mathbf{x}_0, \theta_0) \prod_{l=1}^T p(\mathbf{x}_l, \theta_l | \mathbf{x}_{l-1}, \theta_{l-1}) p(\mathbf{y}_l | \mathbf{x}_l, \theta_l). \end{aligned} \quad (3.9)$$

One of the reasons that the assumption of Gaussian error is common is that if the prior $p(\mathbf{x}_{0:T}, \theta_{0:T})$ or likelihood $p(\mathbf{y}_{1:T} | \mathbf{x}_{0:T}, \theta_{0:T})$ are Gaussian then it follows that the posterior $p(\mathbf{x}_T, \theta_T | \mathbf{y}_T)$ will also be Gaussian, as the product in Equation (3.9) will allow additive powers of the Gaussian exponential. This may significantly simplify the formulation of the posterior, as its form will be known up to a proportionality constant (Vetra-Carvalho et al. 2018).

3.4 Parameter estimation methods

The most common approach to DA for joint-state parameter estimation is to use a state-estimation method but append any parameters θ_n to the state vector \mathbf{x}_n , add the parameter update equation to \mathcal{M}_n in Equation (3.1) (which may be the identity if θ_n is constant) and include state-parameter cross covariance terms in the covariance matrix. Indeed, the most simplistic EnKF formulation for joint state-parameter estimation uses

this method. If the system is highly non-linear however, augmented EnKF for state-parameter estimation may seriously diverge, as the cross covariance between state and parameters may grow exponentially. This is particularly the case for stochastic parameters and when observation frequency is low. Additionally, augmenting the state vector with a large number of parameters may significantly increase the computational load of an EnKF scheme (Santitissadeekorn & Jones 2015).

Various EnKF-based methods to overcome these issues have been formulated. The interacting Kalman Filter, for example, instead implements two separate Kalman Filters to update the state and parameters individually, with the two filters interacting so that state and parameter updates have a physically meaningful association (Koyama & Watanabe 2010). Another option is a parameter augmented version of the local Ensemble Transform Kalman Filter (LETKF), which calculates the Kalman update equation for local regions of the spatial domain, reducing the dimension of calculations and increasing the algorithm's efficiency (Belsky et al. 2013). An alternative is to avoid the augmentation method entirely, thus bypassing the potentially erroneous assumption of a Gaussian joint state-parameter model. Such methods instead estimate parameters based on approximate likelihood functions calculated from the output of a state-only EnKF scheme. This approximation may be done using either a maximum likelihood approach or sequential Bayesian methods (Katzfuss et al. 2016).

Other common Bayesian approaches for non-linear system parameter estimation include Markov Chain Monte Carlo (MCMC) methods and particle filter (PF) methods. MCMC can be inefficient in high-dimensional geophysical models, however, as it requires thousands of model evaluations to calculate the update. Several studies have approached this issue via parallelisation (Solonen et al. 2012) and selective model simplifications (Urban & Keller 2010). Particle filters may also be inefficient in high-dimensional systems, but are useful in that, unlike Kalman Filters, they do not require the assumption of Gaussian error and tend to perform better for stochastic parameters (Santitissadeekorn & Jones 2015). Indeed, PFs are often used as a comparison for testing new DA methods as they can be proven to reproduce the true target posterior distribution in the limit of large particle populations. This desirable property is difficult to prove for many other methods (Law et al. 2015).

Hybrid approaches to parameter estimation are becoming increasingly popular to overcome issues associated with highly non-linear models, combining several DA schemes to exploit the advantages of a variety of methods within one algorithm. Such studies include Santitissadeekorn & Jones (2015), whose 'two-stage' method uses a PF to estimate lower dimensional model parameters while updating the higher dimensional system state using an EnKF. Similarly, Lu et al. (2019) perform a joint parameter-state estimation

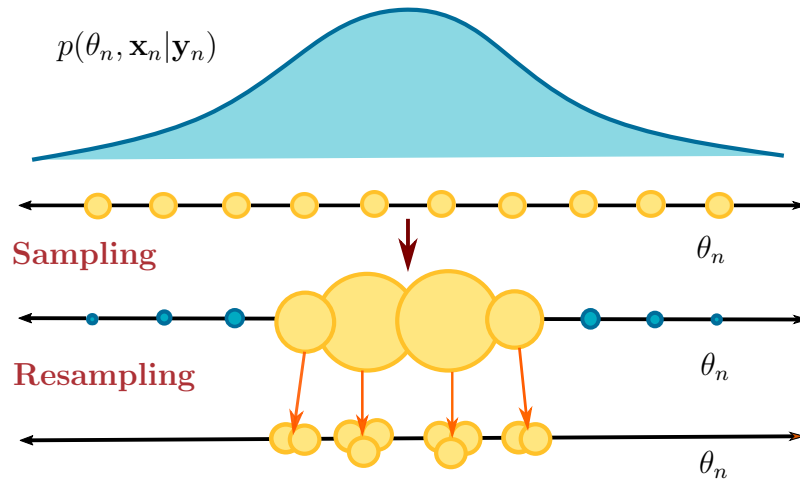


Figure 3.4: Diagram of the process of a simple particle filter. To estimate the posterior, various realisations of the posterior are sampled, referred to as ‘particles’. This can be done by sampling θ_n^i from a prior and then using θ_n^i within the model to compute the corresponding \mathbf{x}_n^i . The particles are weighted according to how closely they match \mathbf{y}_t , the observed data. Heavily weighted particles are then duplicated to retain the same number of members in the particle ensemble. Resampling resets the weights to be uniform across the ensemble, so that particle degeneracy is avoided.

on a stochastic PDE model for paleo-climate by combining an MCMC method with a PF scheme. Such methods, however, require simplifications of the coupling between the various DA schemes used, such as by passing ensemble means of parameters and states between different DA schemes. Hybrid approaches are generally useful if the parameter dimension is much smaller than the state dimension, but parameter estimation for scenarios where the parameter space is high-dimensional still presents significant research challenges.

3.5 Particle Filters

As mentioned, particle filters are a common statistical method for parameter estimation as, although suffering issues in high dimensional systems, they can be shown to converge to the true posterior in the large sample limit. The process of a simple particle filter for

parameter estimation is shown in Figure 3.4, which aims to estimate the target posterior $p(\theta_n, \mathbf{x}_n | \mathbf{y}_n)$ at time t_n . An initial sample of particles is produced from the model by first generating a sample of N parameters θ_n^i from the prior $p(\theta_n)$ and then using these within the model to generate N states $(\mathbf{x}_n, \theta_n^i)$ at time t_n . The initial particles are then weighted according to how closely $h(\mathbf{x}_n, \theta_n^i)$ ‘matches’ \mathbf{y}_n (shown in the top row of Figure 3.4), with the degree of correspondence measured via the innovation $\mathbf{y}_n - h(\mathbf{x}_n, \theta_n^i)$. Heavier weights imply the best matches, while lighter weights imply little correspondence between \mathbf{y}_n and $h(\mathbf{x}_n, \theta_n^i)$. After the weights are calculated, the particle population is updated, so that the greatest number of population members correspond to values of $(\mathbf{x}_n, \theta_n^i)$ with the highest weights. An additional resampling step may also be performed at this point, which, if there is a lack of diversity in the particle population, draws a population of the same size with probabilities proportional to the weight vector. It then resets their weights to be $1/N$, so that the majority of the weight does not become concentrated on a single particle (shown in the bottom row of Figure 3.4) (Vetra-Carvalho et al. 2018, Leeuwen 2015).

Formally, the joint model-parameter posterior $p(\mathbf{x}_l, \theta_l)$, as represented by N realisations of the system state and parameter, $\{\mathbf{x}_l^i, \theta_l^i\}_{i=1}^N$, at time t_l , can be written as

$$p(\mathbf{x}_l, \theta_l) \approx \frac{1}{N} \sum_{i=1}^N \delta(\mathbf{x}_l - \mathbf{x}_l^i, \theta_l - \theta_l^i), \quad (3.10)$$

where $(\mathbf{x}_l, \theta_l) \in \mathbb{R}^{M+1}$ is the model state and $\delta(\mathbf{x} - \hat{\mathbf{x}}, \theta - \hat{\theta})$ is the $M + 1$ dimensional Dirac-delta distribution centred at $\hat{\mathbf{x}}$ and $\hat{\theta}$. This formulation is more generally used in Monte-Carlo methods, and concentrates probability density around ensemble states $(\mathbf{x}_l, \theta_l^i)$ which are closest to (\mathbf{x}_l, θ_l) (Vetra-Carvalho et al. 2018, Madja & Harlim 2012). It follows that Bayes’ Law (3.6) gives the following formulation for the posterior at time t_l ,

$$p(\mathbf{x}_l, \theta_l^i | \mathbf{y}_l) = \frac{p(\mathbf{y}_l | \mathbf{x}_l, \theta_l^i)}{p(\mathbf{y}_l)} p(\mathbf{x}_l, \theta_l^i).$$

Using (3.10) and the Markovian assumption of (3.7), the posterior may be approximated as the weighted sum

$$p(\mathbf{x}_l, \theta_l^i | \mathbf{y}_l) \approx \sum_{i=1}^N w_l^i \delta(\mathbf{x}_l - \mathbf{x}_l^i, \theta_l - \theta_l^i),$$

where the weights are given by

$$w_l^i \propto p(\mathbf{y}_l | \mathbf{x}_l, \theta_l^i) w_{l-1}^i.$$

Each weight w_{l-1}^i is the product of the weights from all previous time steps $t_0 < t_k \leq t_{l-1}$.

The algorithm which arises from this formulation is the simplest of the PF algorithms, known as *Sequential Importance Resampling* or the *Bootstrap PF* (Law et al. 2015). Suppose that at time t_{n-1} we have the posterior distribution represented by the particles $((\mathbf{x}_{n-1}^i, \theta_{n-1}^i), w_{n-1}^i)$, where $(\mathbf{x}_{n-1}^1, \theta_{n-1}^1), \dots, (\mathbf{x}_{n-1}^N, \theta_{n-1}^N)$ is the state ensemble and $w_{n-1}^1, \dots, w_{n-1}^N$ are the N particle weights. These weights must satisfy $w_{n-1}^i \geq 0, \forall i$ and $\sum_{i=1}^N w_{n-1}^i = 1$. Using this set-up, the core steps of the Bootstrap PF algorithm are shown in Algorithm 1 (Reich & Cotter 2015).

Algorithm 1 The Bootstrap PF/ Sequential Importance Resampling

for n in $1 : J$ **do**

1. **Prediction step:** Propagate the ensemble of N states forward from time t_{n-1} to time t_n , $(\mathbf{x}_{n-1}^i, \theta_{n-1}^i) \mapsto (\mathbf{x}_n^i, \theta_n^i)$. The forecast probability distribution then consists of:
 - an ensemble of states $\{\mathbf{x}_n^i, \theta_n^i\}_{i=1}^N$ and
 - a corresponding collection of weights $\{w_{n-1}^i\}_{i=1}^N$.
2. **Filtering step:** Update the weights $\{w_{n-1}^i\}_{i=1}^N$ using the observation \mathbf{y}_n . In the standard PF $w_n^i = c_n w_{n-1}^i p(\mathbf{y}_n | \mathbf{x}_n^i, \theta_n^i)$, where c_n is a normalising constant such that $\sum_{i=1}^N w_n^i = 1$.

end for

Despite its effectiveness in non-linear applications, this standard PF algorithm suffers from several issues. If the system is high-dimensional, PFs suffer from a phenomenon known as particle degeneracy, where one of the particle weights approaches 1, while all others approach zero. When this occurs the approximate posterior effectively has zero width and future steps of the PF will tend to conserve the degenerate weighting. This issue is commonly dealt with by monitoring the Effective Sample Size (ESS) and resampling when the ESS drops below some threshold to refresh the particle ensemble. The resampling step then draws a new ensemble which has the same mean and variance as the old, and sets the weight of each member to be $\frac{1}{N}$ when it is detected that particle weights are becoming degenerate. The greater tendency of PFs towards degeneracy for high-dimensional systems is often referred to as *the curse of dimensionality*, a phenomenon suffered by all importance sampling algorithms in which efficiency decreases rapidly with increasing dimension of the state space (Surace et al. 2019). In high-dimensional spaces, N must be large to give an appropriate estimate of the posterior, which severely decreases the computational efficiency. In fact, it has been shown that the theoretical required sample size to avoid particle degeneracy scales exponentially with the state dimension, rendering PFs very computationally inefficient for high-dimensional systems (Snyder et al. 2008).

3.6 ABC Algorithms

Perhaps one of the simplest Bayesian techniques for parameter estimation is an Approximate Bayesian Computation (ABC) algorithm. ABC may provide an advantage over other Bayesian methods for parameter estimation in that it can be used to compute the posterior distribution, $p(\mathbf{x}_n, \theta_n | \mathbf{y}_n)$, when the likelihood $p(\mathbf{y}_n | \mathbf{x}_n, \theta_n)$ is unknown or intractable. Rather than a likelihood, ABC only requires the ability to sample data from the likelihood at a time t_n . This sample data from the likelihood may be generated by collecting realisations of the parameter θ_n^i from a prior distribution $p(\theta_n)$, and then generating corresponding $\mathbf{x}_n^i \sim p(\mathbf{x}_n | \theta_n^i)$ samples. Each realisation θ_n^i is then accepted if

$$\rho(\mathbf{y}_n, h(\mathbf{x}_n^i)) \leq \kappa, \quad (3.11)$$

where \mathbf{x}_n^i is the i th realisation of the model update for the system at time t_n , ρ is a distance function and κ is a fixed tolerance. The accepted values of θ_n^i will then be samples from the joint posterior $p(\theta_n, \mathbf{x}_n | \rho(\mathbf{y}_n, h(\mathbf{x}_n)) \leq \kappa)$. If κ is sufficiently small then $p(\theta_n, \mathbf{x}_n | \rho(\mathbf{y}_n, h(\mathbf{x}_n)) \leq \kappa)$ should be a reasonable approximation to the target posterior $p(\mathbf{x}_n, \theta_n | \mathbf{y}_n)$. Thus, a likelihood free route to the target posterior is achieved (Sisson et al. 2007).

An illustration of how this process operates is shown in Figure 3.5. We begin with some prior distribution for θ_n based on its physical meaning, in this case assumed to be Gaussian. Many realisations θ_n^i are then drawn from this initial distribution, three of which are shown in Figure 3.5. For each realisation, a forecast \mathbf{x}_n^i is created by running the model forward to the current time using θ_n^i , and then transforming \mathbf{x}_n^i to be in the same space as the observations using the observation operator $h(\cdot)$. Each forecast is then compared to collected observational data \mathbf{y}_n using a comparison which indicates how well \mathbf{x}_n^i ‘matches’ \mathbf{y}_n . Based on this comparison the distribution of the parameter θ_n is updated by selecting only the best matching θ_n^i . These steps are repeated iteratively to further refine the parameter distribution.

Pseudo-code for this ABC algorithm is shown in Algorithm 2 (Gelman et al. 2013). An alternative approach to this algorithm is to select the tolerance κ such that only the ‘best’ $\alpha\%$ of realisations according to the distance function ρ are accepted. For example, selecting the θ_n^i s corresponding to the smallest 20% of distances. Instead of line 6 in Algorithm 2, all θ_n^i are retained and then compared on line 10, with $\alpha\%$ of the best θ_n^i s according to distance being retained to form the updated prior. The particle population can then be replenished by replicating the new population of θ_n^i s to make up the original population size and adding a small amount of zero mean Gaussian noise to the replicates.

Although computationally expensive in its naive form, alterations to the basic ABC algorithm have allowed ABC to be used to estimate parameters associated with complex

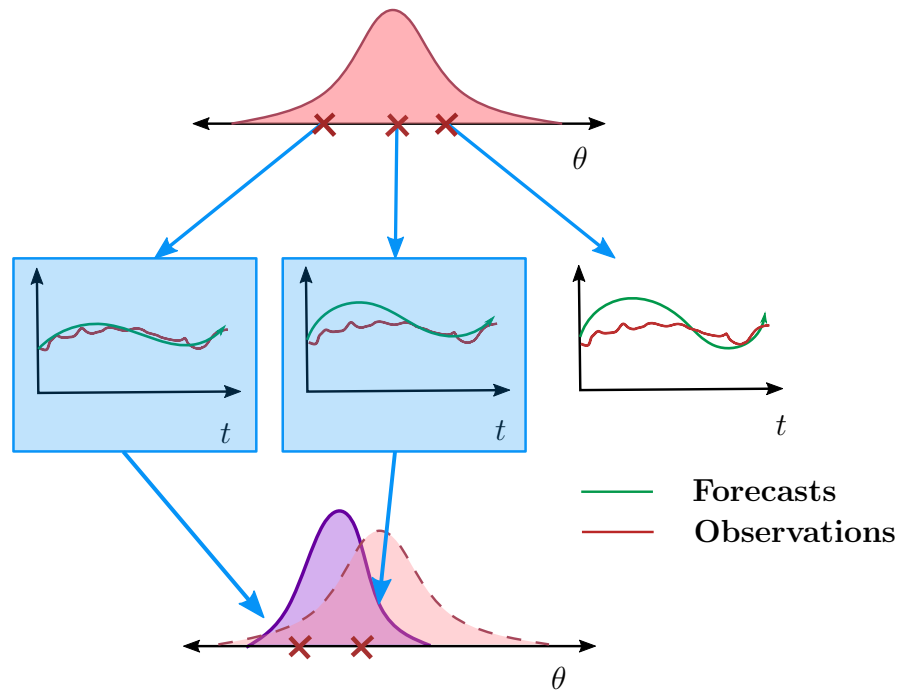


Figure 3.5: Schematic of one assimilation step of an Approximate Bayesian Computation (ABC) type DA scheme. Multiple realisations of the parameter θ_n are drawn from the prior distribution and used within the model \mathcal{M}_n to compute forecasts \mathbf{x}_n^i , for each θ_n^i . These forecasts are compared to collected observational data \mathbf{y}_n via a distance measure ρ , such as an L^2 -norm. Forecasts and observations which most closely match according to the chosen measure are retained and their associated values of θ_n^i used to refine the assumed distribution of θ_n . This is repeated at each time step.

Algorithm 2 The ABC algorithm

- 1: **for** n in $1 : J$ **do**
 - 2: At time t_n collect observational data \mathbf{y}_n .
 - 3: **for** i in $1 : N$ **do**
 - 4: Sample θ_n^i from the prior distribution $p(\theta_n)$.
 - 5: Sample \mathbf{x}_n from the likelihood $p(\mathbf{x}_n | \theta_n^i)$.
 - 6: **if** $\rho(\mathbf{y}_n, h(\mathbf{x}_n^i)) \leq \kappa$ **then**,
 - 7: Retain θ_n^i .
 - 8: **end if**
 - 9: **end for**
 - 10: Update prior distribution $p(\theta_{n+1})$ to be for θ_n^i such that $\rho(\mathbf{y}_n, h(\mathbf{x}_n^i)) \leq \kappa$.
 - 11: **end for**
-

dynamical systems. Tina et al. (2008), for example, develop an ABC type algorithm based on Sequential Monte Carlo methods (SMC), which successfully estimates parameters from biological experimental data. Similarly, Scranton et al. (2014) use SMC-ABC to estimate parameters associated with a stage-structured population model. For such studies the data is complex and the model multi-layered, such that a likelihood free approach is beneficial despite the potential pitfalls of ABC.

3.7 Lagrangian Data Assimilation

Lagrangian data is data which carries information about the trajectory of a tracked object. This data is very common in oceanographic studies, where instruments such as surface and subsurface drifters, gliders and autonomous underwater vehicles are increasingly becoming indispensable in data collection (Apte & Jones 2013). Such data contain detailed information about flow dynamics, which is not captured in classic Eulerian data. As discussed in Chapter 2, Lagrangian data carries not just tracer trajectory information, but detailed information about evolving time-dependent structures in the flow, such as eddies, vortices and gyres. This structural information illuminates barriers and transport mechanisms of key physical quantities, such as energy, and may crucially impact mass flow behaviour.

Lagrangian Data Assimilation (LaDA) is a framework of augmented DA schemes designed to take advantage of the informational richness of Lagrangian data to estimate a dynamical system's state or associated parameters. Ide et. al. (2002) were the first to formulate a formal framework for performing data assimilation on Lagrangian data. They achieved this by augmenting the model governing equation with equations describing tracer advection and using these equations to track the evolving correlation between tracers and the original flow state. This classic Lagrangian data assimilation set-up defines the system state vector \mathbf{x}_t at time t as a vector comprised of both the original flow state \mathbf{x}_t^F and Lagrangian drifters \mathbf{x}_t^D such that

$$\mathbf{x} := \begin{bmatrix} \mathbf{x}_t^F \\ \mathbf{x}_t^D \end{bmatrix}.$$

The state \mathbf{x}_t^F is some property of the fluid, such as temperature, position or velocity, and can be thought of as the Eulerian state vector, which evolves through time according to Eulerian governing equations. Comparatively, the tracer vector \mathbf{x}_t^D is a set of Lagrangian drifters following distinct Lagrangian trajectories. Essentially, the former contains the evolving overall system state, while the latter tracks a state as it evolves along a particular Lagrangian trajectory. The key difference in \mathbf{x}_t^F and \mathbf{x}_t^D is that \mathbf{x}_t^F 's evolution through

time depends only on \mathbf{x}^F at previous times, such that

$$\mathbf{x}_t^F = \mathcal{M}_F(\mathbf{x}_{t_0}^F, t),$$

is the evolution equation for the flow state. Comparatively, the instantaneous tracer position \mathbf{x}_t^D at a given time t gives no information about the flow state \mathbf{x}_t^F , although its evolution equation is dependent on \mathbf{x}_t^F :

$$\mathbf{x}_t^D = \mathcal{M}_D(\mathbf{x}_{t_0}^D, \mathbf{x}_{t_0}^F, t).$$

Hence, LaDA in its most general form may be thought of as DA where the state vector tracks not just model states but also Lagrangian tracer positions, with consequential augmentations to how the state and tracers are evolved through time (Kuznetsov et al. 2003).

The development of LaDA has been challenged predominantly by the complexity of Lagrangian trajectories, which tend to not conform to various assumptions that most non-Lagrangian assimilation schemes rely on (Apte et al. 2008). For example, most DA schemes rely on variables being computed on a fixed grid, whereas Lagrangian observations are dispersed non-uniformly in space and often are not formulated in terms of model variables (Kuznetsov et al. 2003). Many DA schemes also assume Gaussian noise and zero correlation in accumulated error for observed drifter positions, which cannot be assumed to hold generally for Lagrangian data (Kuznetsov et al. 2003, Salman et al. 2006). Additionally, the high order of the non-linearities associated with Lagrangian data may significantly impact the prediction error of many DA schemes formulated for non-Lagrangian non-linear models, like the EnKF (Apte & Jones 2013).

Ide et. al. revealed some of these issues in their seminal paper (2002), finding that with low-quality data and large time intervals between observations, their Lagrangian extended Kalman Filter scheme diverged. They attributed this to the non-linear effects generated by the exponential divergence of tracer trajectories when in the neighbourhood of hyperbolic points of the flow velocity field. Later work by Spiller et al. (2008) approached this issue with a Lagrangian assimilation scheme based on an ensemble Kalman Filter, which allowed issues associated with linearly approximating the error covariance matrix, that tracks correlations between drifter positions and state variables, to be ameliorated. More recently, Apte et al. (2008) developed a Bayesian approach to the issue of non-linearly evolving correlations by using Markov Chain Monte Carlo to sample the ensemble posterior. Other interesting approaches include the ‘‘hybrid filtering’’ methods developed by Slivinski et al. (2015), which use a similar approach to methods developed for parameter estimation in non-linear systems. Similarly to Santitissadeekorn & Jones (2015), they use an EnKF to update high-dimensional state variables and a PF to update lower dimensional drifter positions. Such methods, however, often rely on low dimensional drifters and LaDA methods which are able to deal with high dimensional drifter systems

are an active area of research. Such research aims to exploit the informativeness of Lagrangian data while managing the intertwined issues of non-linearity and error divergence.

3.8 Lagrangian Coherent Data Assimilation

Lagrangian Coherent Data Assimilation (LaCoDA) is a very recent area of LaDA which attempts to tackle the challenges of assimilating Lagrangian Data while also exploiting its inherent structural information through Lagrangian Coherent Structure Theory. This area of research has shown success in tackling highly chaotic flow regimes, as well as flow data which contains significant small scale structures, such as ocean data with sub-meso scale eddies. A study by Maclean et. al. (2017), for example, estimates parameters from Lagrangian ocean drifter data using Principle Component Analysis (PCA) to identify relevant LCS. Their study shows that, when the model exhibits chaotic advection, the incorporation of PCA into an ABC scheme yields improved error in estimating model parameters compared to a Bootstrap PF. This suggests that the incorporation of LCS into LaDA has promise for improving data assimilation error in chaotic regimes.

Studies by Gaultier et. al. (2013, 2014) have similarly attempted to integrate LCS information into LaDA, using the FTLE to extract structural information from Sea Surface Temperature gradients. The studies then use this structural information within Bayesian DA schemes to correct predictions of Sea Surface Height. The latter study uses a high-resolution numerical model of the Soloman Sea, while the former performs the same analysis on altimetric data from the Western Mediterranean Sea. Both data sets exhibit high variability on both submeso- and meso-scales, suggesting the potential for using LCS in LaDA to better utilise fine-scale structural information.

Other studies have used heuristic clustering methods derived from machine learning techniques to group regions of the flow according to coherent structures. Husic, Schleuter-Cluck and Dabiri (2019) develop a Simultaneous Coherent Structure Colouring (SCSC) method which clusters data so that the most dissimilar data points are in separate clusters. Schleuter-Cluck and Dabiri (2019) then use SCSC within parameter estimation schemes, finding that using their coherent colouring method provides a greater level of robustness and accuracy than just using particle displacement alone.

Each of these studies suggest a meaningful advantage in consulting LCS techniques when assimilating challenging Lagrangian data, particularly that exhibiting chaotic behaviour or complex fine scale structure. This motivates the main question postulated in this thesis: is there a performance advantage to using an FTLE-based data assimilation scheme for parameter estimation in chaotic systems? An FTLE-based scheme builds on

the current literature of LaCoDA as the FTLE is designed to tackle spatially traversing coherent structures, unlike the PCA method in Maclean et al. (2017). Such moving structures are a common feature of geophysical flows and particularly in Lagrangian data, so methods which can exploit such structures are desirable. Although the FTLE had been used previously to aid geophysical predictions in Gaultier et. al. (2013, 2014), they use the FTLE as a means of correcting errors in velocity field data, rather than as a variable, and do not consider multiple time-steps of an assimilation scheme. Additionally, these studies only incorporate information about the positions of FTLE ridges, a process which is itself ambiguous in finite-time flows, as discussed in Section 2.5. Comparatively, the investigation herein will consider the use of the FTLE for parameter estimation over multiple time steps of a DA scheme and using the entire FTLE field. Lastly, the FTLE is a non-linear technique and hence may accommodate the non-linearity challenges which thwart many LaDA schemes. Although this non-linearity makes the formulation of key DA objects such as likelihoods challenging, the use of an ABC scheme may sidestep this issue, as it only requires the ability to sample from the likelihood. These considerations motivate Chapter 4, which develops the main premise of this thesis: an FTLE-based ABC scheme for parameter estimation in chaotic regimes of Rossby wave flow.

Chapter 4

FTLE-ABC

With the motivation of improving parameter estimation in highly chaotic geophysical flows, Chapter 4 develops an FTLE-based ABC scheme for estimating the ϵ parameter in chaotic regimes of Rossby wave flow. Following a formulation of the FTLE-ABC algorithm, the predominant research questions associated with testing the algorithm's efficacy are discussed. The algorithm's performance is then qualitatively compared to the benchmark algorithms of standard ABC and a standard particle filter in a variety of scenarios, including for different values of the true parameter ϵ , varied assimilation time-steps, different initialisation schemes, and model reference frames.

4.1 The FTLE-ABC Algorithm

The method developed here, referred to as FTLE-ABC, is analogous to the ABC-type assimilation scheme for parameter estimation first described in Figure 3.5 and Algorithm 2. Figure 4.1 shows the augmented scheme for FTLE-ABC. The assimilation process begins identically to the standard ABC scheme, by drawing multiple realisations from some chosen prior distribution for the parameter being estimated. As in the case of Figure 3.5, the target posterior is that for the parameter being estimated, here the Rossby wave flow parameter ϵ . However, rather than simply using the positions of fluid particles at the j th time, t_j , their FTLE field over the time period $[t_{j-1}, t_j]$ is used.

Figure 4.1 illustrates in detail how the FTLE-ABC scheme proceeds numerically. After realisations of the parameter ϵ are drawn from a prior distribution (shown here as uniform), for each realisation ϵ^i the system is propagated forward in time using the model \mathcal{M}_n and the gridding method explained in Figure 2.3. The gridding method allows a flow map gradient to be estimated and from this a forecast FTLE field, $FTLE(\mathbf{x}_j^i)$, is calculated for the model state \mathbf{x}_j^i corresponding to each realisation ϵ^i . Using the same

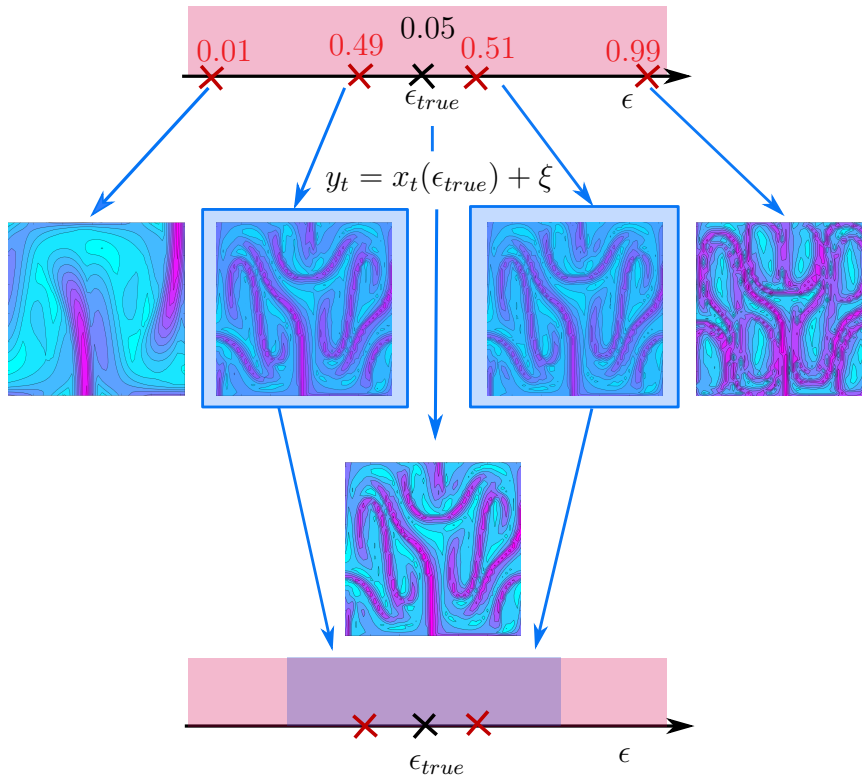


Figure 4.1: Schematic of one assimilation step of the FTLE-ABC DA scheme. Realisations of the parameter ϵ are drawn from a prior distribution (in this case uniform). For each realisation the system is propagated forward in time using the gridding method explained in Figure 2.3. This is used to calculate the flow map gradient and hence an FTLE field for each parameter realisation, which will form the forecast FTLE. An observational FTLE field is also calculated from tracer data collected at the same time. The forecast and observed FTLE fields are compared via an L^2 -norm distance measure. The best $\alpha\%$ of these ϵ values according to the distance measure are retained to form the prior for the next assimilation step .

method, an observed FTLE field, $FTLE(\mathbf{y}_j)$, is also calculated from tracer data \mathbf{y}_j . If using real data for \mathbf{y}_j , corresponding velocities can be estimated by using a suitable finite difference method, so that the observed flow map gradient can be estimated and used to calculate an observed FTLE field. If data is being simulated, as was done in the investigations in Chapters 4 and 5, the gridding method can be used with initial particle positions calculated from the model with the true parameter as input and a small amount of Gaussian noise added to simulate data error. The forecast and observed FTLE fields are then compared via a distance measure ρ from the condition (3.11), which we choose to be the L^2 -norm here. The forecast FTLE fields most similar to the observed FTLE field according to this distance measure, and their associated ϵ^i values, are used to further refine the target posterior distribution for ϵ . From Figure 4.1 it can be appreciated how forecast FTLE fields corresponding to similar ϵ values appear structurally similar, while those with strongly differing ϵ values bear little structural resemblance. A higher distance between a particular forecast FTLE field and the observed FTLE field should imply the structures present in each are poorly matched, suggesting ϵ^i differs strongly from ϵ_{true} . A lower distance should similarly imply a good match between the structures in the observed FTLE field and forecast FTLE field, suggesting the corresponding ϵ^i is closer to ϵ_{true} .

Algorithm 3 The FTLE-ABC Algorithm

```

for  $j = 1 : J$ , do
  At time  $t_j$  collect tracer observations  $\mathbf{y}_j$ .
  Use the observations  $\mathbf{y}_j$  to calculate the forecast FTLE field,  $FTLE(\mathbf{y}_j)$ .
  for  $i$  in  $1 : N$ , do
    Sample  $\epsilon_j^i$  from the prior distribution  $p(\epsilon_j)$ .
    Sample  $\mathbf{x}_j^i$  from the likelihood  $p(\mathbf{x}_j^i, \epsilon_j^i | \mathbf{y}_j)$  by forwarding the model to time  $t_j$ 
    for  $\epsilon_j^i$ .
    Calculate the observed FTLE field for  $\epsilon_j^i$ ,  $FTLE(\mathbf{x}_j^i)$ .
    if  $\|FTLE(\mathbf{y}_j) - FTLE(\mathbf{x}_j^i)\|_2 < \kappa$  then
      Retain  $\epsilon_j^i$ 
    end if
  end for
  Update prior  $p(\epsilon_{j+1})$  to be retained  $\epsilon_j^i$  only.
  Replenish population size by creating replicates with added non-zero Gaussian noise.
end for

```

The relative magnitudes of the L^2 -norm for each parameter realisation are used to update the target posterior here by retaining the best matched ϵ^i and discarding the worst. This is done by selecting the best $\alpha\%$ of ϵ^i according to distance and replenishing the particle population by generating a sufficient number of replicates of this best $\alpha\%$ and adding Gaussian noise to avoid degeneracy. For example, for $N = 100$ and $\alpha = 20$, the

ϵ realisations corresponding to the smallest 20 distances are retained. To replenish the population to size 100, 4 copies of the 20 selected ϵ are generated, with a small amount of zero-mean Gaussian noise added to each. This FTLE-ABC algorithm is given in pseudo-code in Algorithm 3. Note that in Algorithm 3 the value κ is chosen to be the minimum L^2 -norm distance not included in the best $\alpha\%$ of particles.

4.2 Research questions for FTLE-ABC

The investigation herein examines the performance of the FTLE-ABC algorithm when estimating the ϵ parameter in Rossby wave flow, with the predominant aim of exploring the algorithm's efficacy in chaotic flow regimes. In Section 2.8 connections between the flow chaoticity and the ϵ parameter of Rossby wave flow, as well as T_f , the time over which the assimilation scheme is performed, were discussed. There it was emphasised how certain values of ϵ tend to maximise the chaoticity of the flow for a particular T_f . The time frame over which chaotic behaviour occurs, the chaotic timescale, may be approximated as $1/\text{FTLE}_{med}$, and will impact how effectively the chosen assimilation time step, ΔT , captures the system dynamics. As the predominant aim of this thesis is to develop a parameter estimation method which is resilient in chaotic geophysical systems, the following numerical investigations focus on comparing the performance of FTLE-ABC, standard ABC and a standard particle filter for different magnitudes of the ϵ and ΔT . As ϵ scales the degree of chaos in the flow, and ΔT affects how well the scheme captures the time resolution of the chaotic dynamics, varying these parameters will give insight into how FTLE-ABC's performance is impacted by chaotic behaviour.

Finding DA schemes which perform well for large gaps between data collection times is also of interest as this reduces the need for highly frequent (and hence more expensive) data collection and storage. As discussed in Chapter 1, data with large times between collection periods contains noise which is magnified by chaotic behaviour, particularly if the model is run over many assimilation time steps for large final times. It is consequently of interest whether FTLE-ABC may gain an advantage over standard methods for longer times between assimilation steps (ΔT), as the technique relies on the formation of coherent structures, rather than just tracer positions. Longer times between assimilation steps allows more time for evolving coherent structures to develop and hence may actually enhance the information captured in an FTLE field, rather than hindering the method's performance. These factors are an additional motivation for analysing the algorithm's performance for different magnitudes of ΔT .

Apart from these primary concerns of dealing with chaoticity and infrequent data collection, it is also of interest to assess FTLE-ABC's ability to deal with spatially traversing

coherent structures. Generally, geophysical fluid structures are temporally *and* spatially dynamic, so LaCoDA methods which can deal with spatially dynamic fluid structures are advantageous. In particular, such a property would improve on Maclean et al. (2017), as the PCA method they use for structure detection is not frame independent and cannot deal with spatially traversing structures. As the FTLE is frame independent, however, the method can easily track spatially traversing structures, such as those occurring in the moving frame of Rossby wave flow. To demonstrate that FTLE-ABC can deal with spatially traversing structures, the algorithm's performance is compared in the moving and co-moving frames of Rossby wave flow here, as the moving frame case contains spatially traversing structures while the co-moving frame case does not.

An additional concern is whether the initial placement of tracers will affect FTLE-ABC's performance. This issue is relevant to the problem of optimising methods of collecting oceanographic data, where instruments such as temperature sensors may be dropped by boats for data collection. Optimal placement of tracers may impact the performance of LCS based methods and is a question of importance in geophysical applications (Salman et al. 2008). More generally, as is discussed in Chapter 1, geophysical data may be erratically distributed and so gaining insight into how a new method deals with non-uniform data is important. Here this is tested by comparing the assimilation scheme results when all tracers are initially dropped inside the gyres in the Rossby wave model to when initialisation is on a uniform grid. The former approach does not allow uniform sampling of the initial domain, which may impact the formation of FTLE ridges and hence the ability of the FTLE to identify significant structures. Additionally, the gyres of Rossby wave flow tend to be less chaotic, more coherent regions of the flow and so initialisation in gyres only may cause the method to underestimate ϵ , the magnitude of which scales with the chaoticity. Knowledge of the algorithm's sensitivity to the initialisation scheme used will give insight into suitable applications for the algorithm.

These research questions are qualitatively investigated in Sections 4.3 and 4.4. Section 4.3 qualitatively compares the performance of FTLE-ABC and standard ABC by considering individual assimilation steps of the two schemes for different magnitudes of ΔT and ϵ , the two initialisation schemes and the model in the co-moving and moving frames. Section 4.4 compares FTLE-ABC and a standard particle filter, focusing on regimes in which the particle filter becomes degenerate, such as for large ΔT or for high initial noise in tracer positions. This different approach for the PF comparison, which allows insight into how well FTLE-ABC deals with noisy tracer data, is discussed in further detail in Section 4.4.

4.3 Standard ABC versus FTLE-ABC

As an initial investigation into the impact of parameters such as ΔT and ϵ_{true} on the performance of FTLE-ABC and standard ABC, plots of the L^2 -norm distance against ϵ are generated for varied magnitudes of these parameters. For FTLE-ABC this means the value of the distance $\|FTLE(\mathbf{y}_j) - FTLE(\mathbf{x}_j^i)\|_2$ is plotted against its corresponding ϵ_j^i realisation at the j th time step of the algorithm. For standard ABC, the distance $\|\mathbf{y}_j - \mathbf{x}_j^i\|_2$ is plotted against its corresponding ϵ_j^i realisation. The algorithms are run for a larger number of observations ($J = 9$) and then a small number of observations ($J = 3$), so that their performance is compared for the small $\Delta T \approx 1.3$ in the former case, and the large $\Delta T = 4$ in the latter case. These ‘large’ and ‘small’ definitions are designated by the $1/FTLE_{med}$ curve for $T_f = 12$ in Section 2.8, which suggests $\Delta T \approx 1.33$ will be smaller than the chaotic timescale for all ϵ , resolving the chaotic dynamics well, and $\Delta T = 4$ will be larger than or of similar magnitude to the chaotic timescale for $\epsilon > 0.5$. Different magnitudes of ϵ_{true} (0.3, 0.9 and 1.3) are also compared, to gauge the effect of varying the degree of chaoticity on the algorithms’ performance. For each of these comparisons, the first and final assimilation steps ($j = 1, J$) for a final time of $T_f = 12$ are plotted, to illustrate how efficiently each algorithm updates the distribution of ϵ as the assimilation scheme is run. If a method performs well, we expect a distance minimum to form near the true parameter in the first step of the assimilation scheme, as this would indicate particles giving the lowest distances correspond to ϵ realisations very near to ϵ_{true} . In the final step of the assimilation scheme, a well performing method should show a narrow distribution with an average near the true parameter value, as this indicates low error in the scheme’s estimate of ϵ_{true} .

Motivated by the research questions in Section 4.2, FTLE-ABC’s performance is also compared for different initial distributions of tracer data and in conditions where coherent structures may spatially traverse the flow domain. Section 4.3.1 details the simplest comparison between the algorithms, where tracers are initialised on a uniform grid, so that there is uniform sampling of the flow domain, and the Rossby wave model is in the co-moving frame, so that structures cannot spatially traverse the flow domain. In Section 4.3.2 a non-uniform initialisation scheme is used, with tracers being dropped randomly in one of the two gyres of Rossby wave flow, again when the model is in the co-moving frame. This provides a comparison to the uniform initialisation case, to investigate how well FTLE-ABC deals with non-uniform sampling of the initial flow domain. Finally, in Section 4.3.3 the algorithms are compared again for the uniform initialisation scheme, but for the model in the moving frame, so that structures may spatially traverse the flow. This allows a comparison to the performance of FTLE-ABC in Section 4.3.1, where the co-moving frame is used, to determine how well the new algorithm deals with spatially traversing structures.

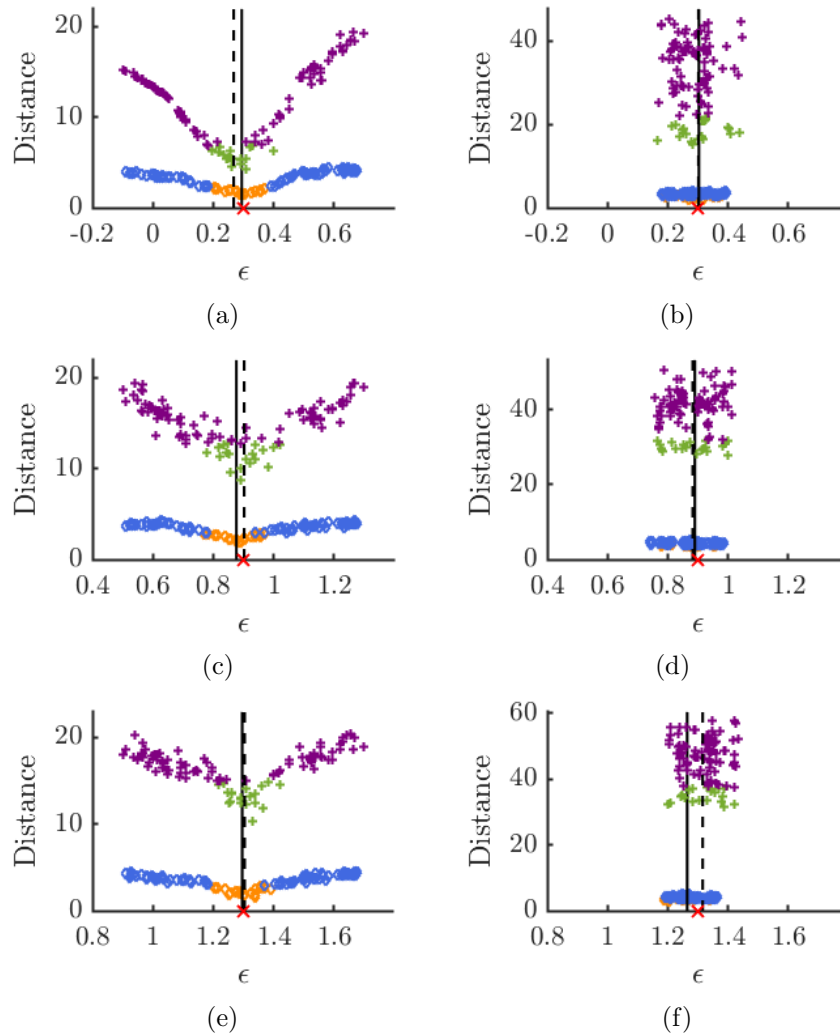


Figure 4.2: Comparison of the first assimilation step (left) and last assimilation step (right) for FTLE-ABC (blue and orange diamonds) and standard ABC (green and purple crosses). The lines represent the average of the best 20% of particles according to distance for FTLE-ABC (black) and standard ABC (black dashed). ϵ_{true} is denoted with a red cross. (a), (b) use $\epsilon = 0.3$, (c), (d) $\epsilon = 0.9$ and (e), (f) $\epsilon = 1.3$. The runs are for the model in the co-moving frame with $T_f = 12$, $J = 9$, $M = 100$, when initialised on a uniform grid.

4.3.1 Spatially stationary structures with uniform initialisation

Here we consider the most simple comparison between standard ABC and FTLE-ABC, where tracers are initialised on a uniform 10-by-10 grid ($M = 100$). The dimension $M = 100$ was chosen following qualitative investigations into the effect of the dimension,

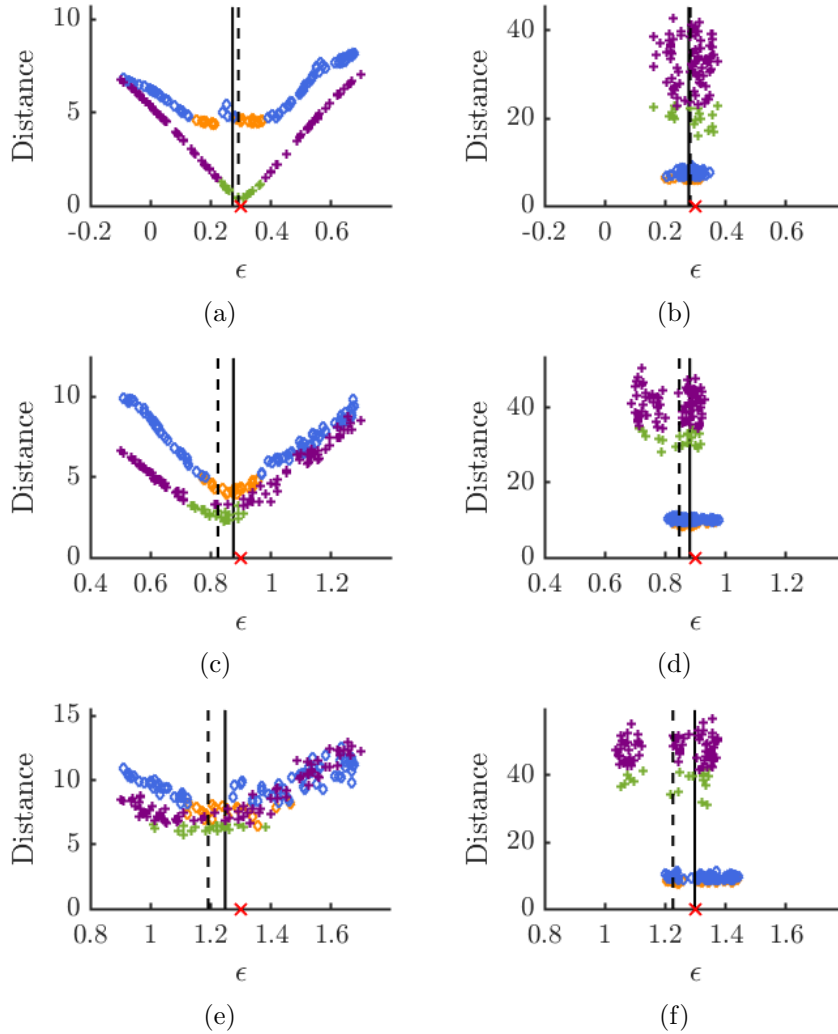


Figure 4.3: Comparison of the first assimilation step (left) and last assimilation step (right) for (a), (b) use $\epsilon = 0.3$, (c), (d) $\epsilon = 0.9$ and (e), (f) $\epsilon = 1.3$, with the same colour scheme as Figure 4.2. The runs are for the model in the co-moving frame with $T_f = 12$, $J = 3$, $M = 100$, when initialised on a uniform grid.

which suggest this is a suitable number of points on the spatial grid to give a suitably high resolution FTLE field, while not being unreasonably computationally time consuming. In this first instance, the model is in the co-moving frame, so that extant structures are spatially stationary. In each of the plots in the following sections, orange and blue diamonds indicate particles from FTLE-ABC runs, while green and purple crosses indicate particles from standard ABC. Orange and green symbols represent the best $\alpha = 20\%$ of particles according to the L^2 -norm distance measure. The black filled line represents the average

value of ϵ for the best 20% of particles in the FTLE-ABC scheme, while the black dashed line represents the average ϵ for the best 20% of particles in the standard ABC scheme. A red cross denotes ϵ_{true} .

Figures 4.2a, 4.2c and 4.2e show the L^2 -norm distance versus ϵ plots for the first assimilation step with $J = 9$ observations, so that the assimilation step is $\Delta T \approx 1.33$. Figure 4.2a is for the relatively low $\epsilon_{true} = 0.3$, Figure 4.2c for $\epsilon_{true} = 0.9$ and Figure 4.2e for the relatively high $\epsilon_{true} = 1.3$. As expected, both schemes give a minimum very close to the true value of ϵ , although the position of the minimum is slightly more diffuse for the higher ϵ_{true} values of 0.9 and 1.3. The average value of the distributions on these plots is also very near the true value of ϵ .

Figures 4.2b, 4.2d and 4.2f show the same three plots but for the final step of the 9 step assimilation scheme. From these plots it can be seen that FTLE-ABC and standard ABC perform similarly for $\epsilon_{true} = 0.3$ and $\epsilon_{true} = 0.9$, giving quite narrow final distributions with an average ϵ which lies on the true value of the parameter. The narrow intervals for these distributions suggests low error in the estimate for both methods. For $\epsilon_{true} = 1.3$ the methods give slightly different final distributions, with standard ABC slightly overestimating ϵ_{true} and FTLE-ABC slightly underestimating ϵ_{true} . The final distribution for FTLE-ABC is also marginally narrower than that for standard ABC, but not significantly so. These plots suggest that for this small assimilation time step when a uniform initialisation scheme is used and only stationary structures are extant in the flow, FTLE-ABC and standard ABC perform equivalently. This seems to hold even for higher chaoticity regimes of the flow (higher ϵ).

Figures 4.3a, 4.3c and 4.3e show the same L^2 -norm distance versus ϵ plots for the first assimilation step, but for a lower number of observations, $J = 3$, so that the assimilation time step is now larger ($\Delta T = 4$). The minimum distance for both methods is again quite near the true parameter for $\epsilon_{true} = 0.3$, suggesting the methods perform similarly in this low chaoticity regime. The initial distributions for the parameters are also again more diffuse for $\epsilon_{true} = 1.3$, spreading between 1 and 1.4, which can be attributed to this being a higher chaoticity regime than for the lower ϵ_{true} values. For the larger ϵ_{true} values of 0.9 and 1.3, standard ABC again tends to underestimate the true parameter, while FTLE-ABC gives a more accurate estimate, with the average of the particle distribution lying near to the true value.

Figures 4.3b, 4.3d and 4.3f show the same plots for the final assimilation step. For $\epsilon_{true} = 0.3$ the final distributions for FTLE-ABC and standard ABC are of a similar width, spreading between 0.2 and 0.35, and have averages which lie on the value of the true parameter, suggesting the methods again perform similarly in this low chaoticity regime. At

the higher ϵ_{true} values, however, the FTLE-ABC final particle distribution becomes narrower than that for standard ABC, particularly for $\epsilon_{true} = 1.3$, giving higher certainty in the ϵ_{true} estimate. For $\epsilon_{true} = 0.9$, standard ABC slightly underestimates ϵ_{true} as around 0.85, while FTLE-ABC estimates the true parameter quite accurately. For $\epsilon_{true} = 1.3$ the difference in the techniques' estimates is more stark, with FTLE-ABC again estimating ϵ_{true} quite accurately, while standard ABC underestimates the true parameter as 1.21. Such differences in ϵ appear small, but have a significant impact on the nature of the flow, as shown in Figure 2.9 in Chapter 2, so small errors in the estimate may have a significant impact on predictions. These plots suggest that for the higher chaoticity regimes of Rossby wave flow, and when there are larger time steps between assimilation steps, FTLE-ABC may more accurately estimate ϵ , and also give less error in this estimation (a narrower final distribution). Thus, for the uniform initialisation and spatially stationary structures case, FTLE-ABC's performance is not severely effected when the time between observations is longer, while standard ABC's performance worsens in the more chaotic regimes induced by longer ΔT and higher ϵ_{true} values.

4.3.2 Spatially stationary structures with non-uniform initialisation

Here the model is again considered in the co-moving frame, so that structures are spatially stationary, however now for the non-uniform initialisation scheme where all tracers are initially placed randomly inside one of the two gyres. This will allow some insight into how the new algorithm deals with non-uniform sampling of the initial flow domain. Figures 4.4a, 4.4c and 4.4e again show L^2 -norm distance versus ϵ plots for the first assimilation step with the lower $\Delta T \approx 1.33$ ($J = 9$). For $\epsilon_{true} = 0.3$ and $\epsilon_{true} = 0.9$ FTLE-ABC and standard ABC again perform similarly, giving minima around the true parameter value. For $\epsilon_{true} = 1.3$ however, the distribution of FTLE-ABC's best 20% of particles tends to be biased towards values of ϵ which are higher than the true value, which seems counter-intuitive given that inside the gyres should be a less chaotic region of the flow domain. Possibly the non-uniform sampling has biased the FTLE field such that the estimate for ϵ is too high. Figures 4.4b, 4.4d and 4.4f show the same plots for the final assimilation step. From these it can be seen that when particles are initially placed randomly in gyres, standard ABC will tend to underestimate ϵ_{true} , particularly for higher ϵ_{true} , possibly due to the gyres being less chaotic regions of the flow. FTLE-ABC overestimates the true parameter for the low $\epsilon_{true} = 0.3$ value, but appears quite accurate for the higher values $\epsilon_{true} = 0.9, 1.3$. This is possibly due to the higher chaoticity allowing more even sampling of the flow domain when the flow has been allowed to evolve for some time, so that FTLE calculations are no longer as biased as in the first assimilation step.

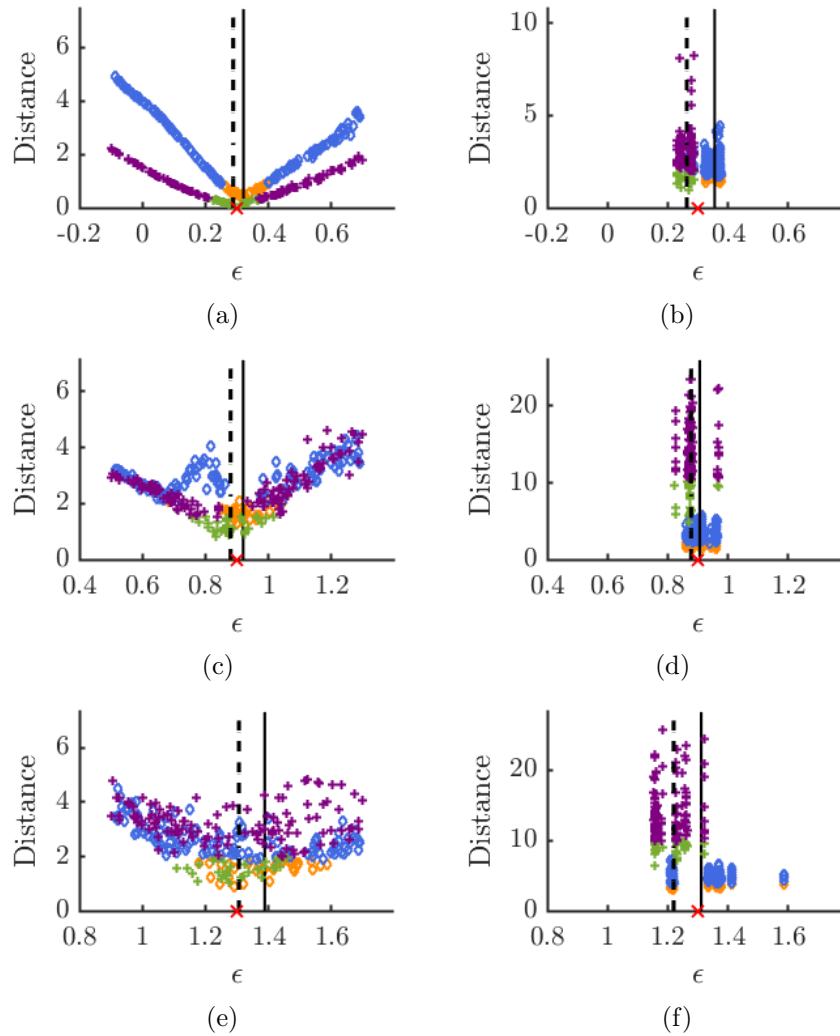


Figure 4.4: Comparison of the first assimilation step (left) and last assimilation step (right) for (a), (b) use $\epsilon = 0.3$, (c), (d) $\epsilon = 0.9$ and (e), (f) $\epsilon = 1.3$, with the same colour scheme as Figure 4.2. The runs are for the model in the co-moving frame with $T_f = 12$, $J = 9$, $M = 100$, when initialised inside gyres.

Figures 4.5a, 4.5c and 4.5e show the same L^2 -norm distance versus ϵ plots for the first assimilation step when initialised randomly within gyres, but for the larger $\Delta T = 4$ ($J = 3$). Interestingly, for the case of $\epsilon_{true} = 0.3$ the distance plot for FTLE-ABC possesses two minima, one at the true value of 0.3 and another at $\epsilon = 0.5$, possibly due to the partitioned initialisation scheme causing the algorithm to assign a separate ϵ_{true} for each of the gyres and creating a bi-modal distribution. For $\epsilon_{true} = 0.9$ only one minimum occurs, while for $\epsilon_{true} = 1.3$ the single minimum is very dispersed. Despite

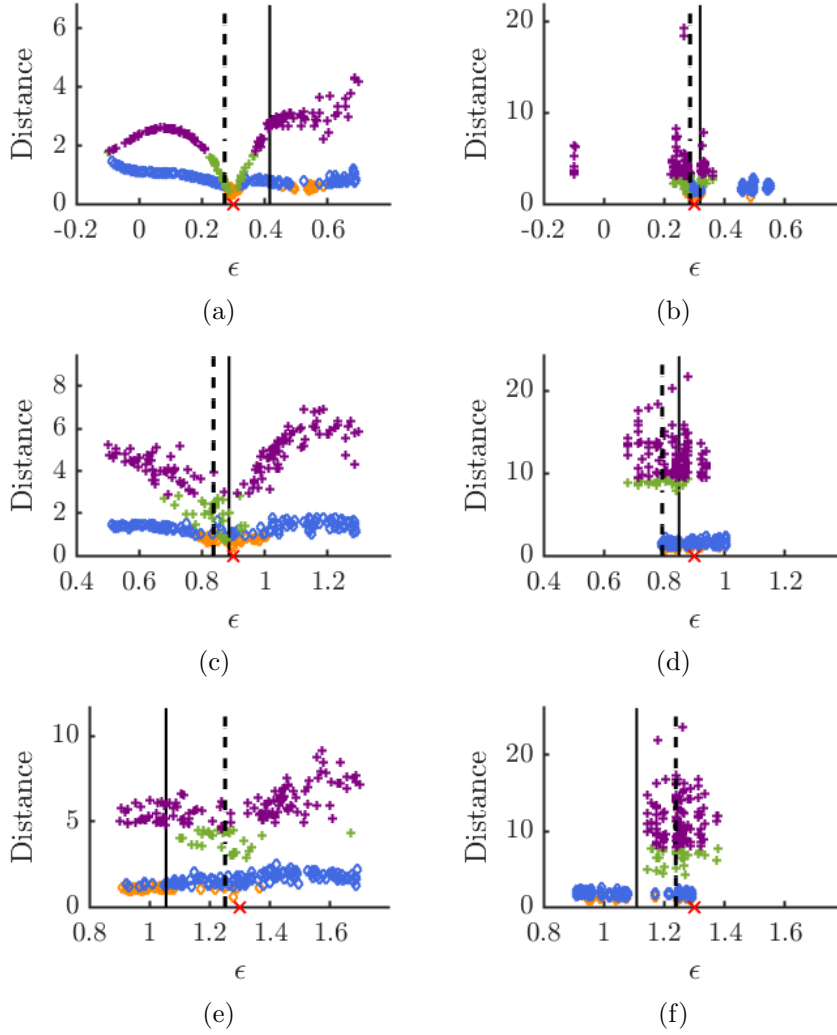


Figure 4.5: Comparison of the first assimilation step (left) and last assimilation step (right) for (a), (b) use $\epsilon = 0.3$, (c), (d) $\epsilon = 0.9$ and (e), (f) $\epsilon = 1.3$, with the same colour scheme as Figure 4.2. The runs are for the model in the co-moving frame with $T_f = 12$, $J = 3$, $M = 100$, when initialised inside gyres.

the distribution becoming more dispersed for higher ϵ_{true} values, for standard ABC a single minima occurs for all ϵ_{true} values and is generally centred on the true parameter value. Figures 4.5b, 4.5d and 4.5f show the same plots for the final assimilation step. Here FTLE-ABC tends to perform badly for higher ϵ_{true} but reasonably well for lower values, which seems to contradict the lower ΔT case. This is possibly due to higher ΔT coupled with the uneven sampling of the flow domain biasing the FTLE field ridges to particular portions of the domain. For example, large ridges may form around the

gyres due to the large ΔT , while other portions of the domain remain empty. Overall these results suggest the FTLE-ABC algorithm produces very mixed results when initial sampling of the flow domain is non-uniform and ΔT is high, possibly because this biases the stretching measures of the FTLE field and creates an incomplete picture of extant coherent structures.

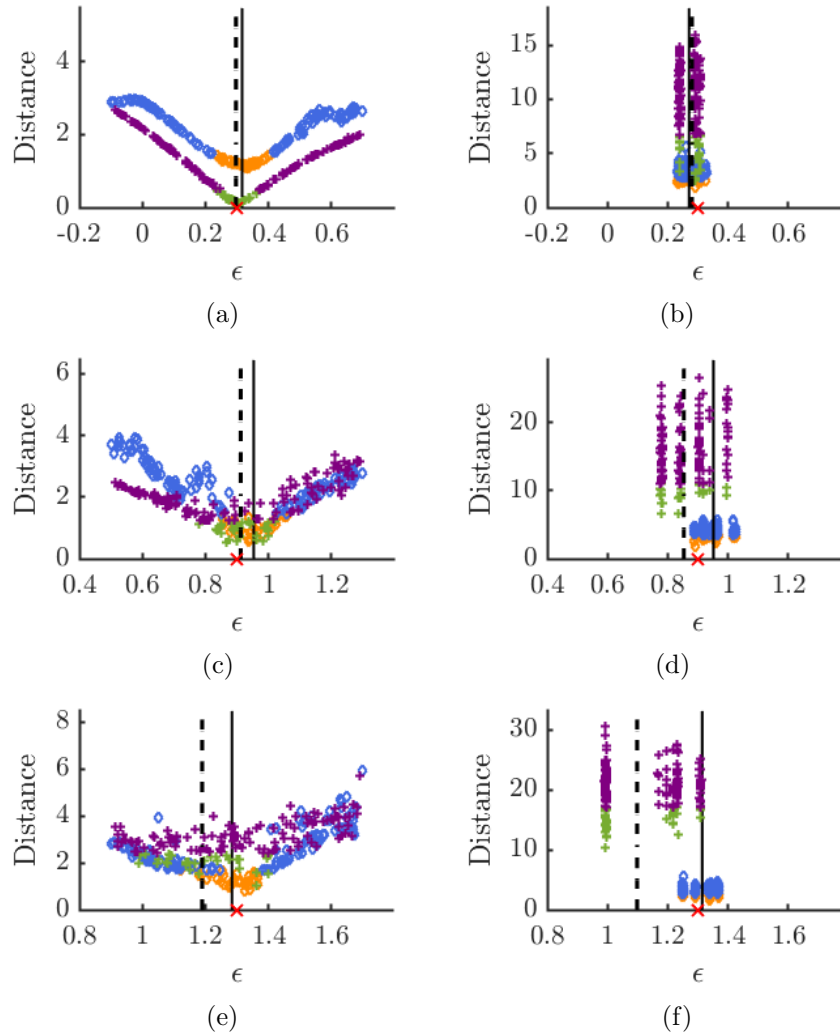


Figure 4.6: Comparison of the first assimilation step (left) and last assimilation step (right) for (a), (b) use $\epsilon = 0.3$, (c), (d) $\epsilon = 0.9$ and (e), (f) $\epsilon = 1.3$, with the same colour scheme as Figure 4.2. The runs are for the model in the moving frame with $T_f = 12$, $J = 9$, $M = 100$, when initialised on a uniform grid.

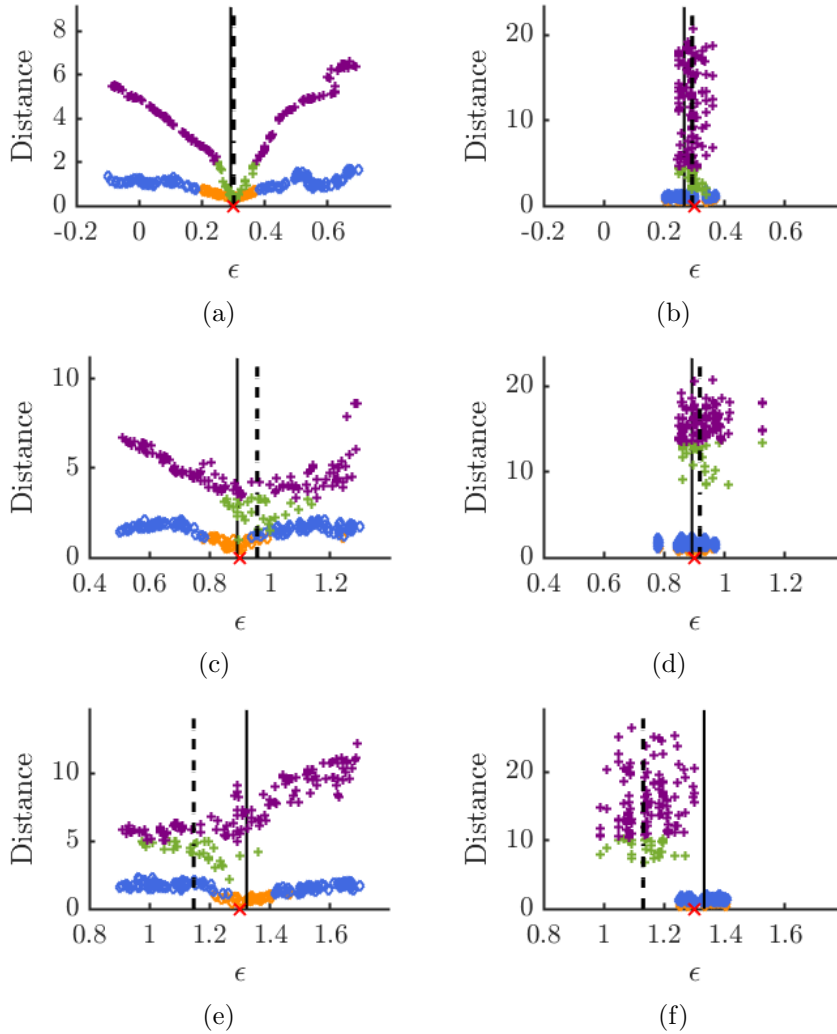


Figure 4.7: Comparison of the first assimilation step (left) and last assimilation step (right) for (a), (b) use $\epsilon = 0.3$, (c), (d) $\epsilon = 0.9$ and (e), (f) $\epsilon = 1.3$, with the same colour scheme as Figure 4.2. The runs are for the model in the moving frame with $T_f = 12$, $J = 3$, $M = 100$, when initialised on a uniform grid.

4.3.3 Spatially traversing structures with uniform initialisation

We now consider FTLE-ABC in the moving frame of Rossby wave flow, where extant structures may spatially traverse the flow domain. Spatially traversing structures are common in chaotic geophysical flows and LCS methods which may deal with such structures are desirable. Figures 4.6a, 4.6c and 4.6e show the L^2 -norm distance versus ϵ plots for the first assimilation step in this scenario, with the uniform grid initialisation and $\Delta T \approx 1.33$. As can be seen, for the first assimilation step, the standard and FTLE-ABC

methods appear to perform similarly for $\epsilon_{true} = 0.3$, forming distinct minima about the true parameter value. For $\epsilon_{true} = 1.3$ the minima for both methods become very diffuse, as expected due to the higher chaoticity, but the average for standard ABC in this case underestimates ϵ_{true} by around 0.1, while FTLE-ABC estimates the true parameter quite accurately. Figures 4.6b, 4.6d and 4.6f show the same plots for the final assimilation step. For $\epsilon_{true} = 0.3$ and $\epsilon_{true} = 0.9$ the methods perform similarly, although FTLE-ABC gives a narrower final distribution for $\epsilon_{true} = 0.9$ than standard ABC. For $\epsilon_{true} = 1.3$, the minimum for standard ABC becomes much more diffuse than that for FTLE-ABC, suggesting the higher chaoticity is causing the error in the algorithm's estimation to grow. In particular, the final distribution ranges from 1.25 to 1.37 for FTLE-ABC but from 1 to 1.3 for standard ABC. The average for the distribution for standard ABC is also around 0.2 lower than the true value here, while FTLE-ABC estimates the true value quite accurately. These results suggest that, as expected, the performance of FTLE-ABC does not change significantly when spatially traversing structures are extant in the flow, while that of standard ABC is quite severely effected by spatially traversing structures when in higher chaoticity flow regimes.

Figures 4.7a, 4.7c and 4.7e show the same L^2 -norm distance versus ϵ plots for the first assimilation step, but now for the higher time step, $\Delta T = 4$. The two methods perform similarly in this first step, except for when $\epsilon_{true} = 1.3$, where standard ABC underestimates the true parameter by about 0.2 in this more chaotic regime. Figures 4.7b, 4.7d and 4.7f show the same plots but for the final assimilation step. Here, both methods perform similarly for $\epsilon_{true} = 0.3$ and $\epsilon_{true} = 0.9$, although standard ABC's distribution for $\epsilon_{true} = 0.9$ is again more diffuse, giving higher variance in the method's estimate for ϵ_{true} . For $\epsilon_{true} = 1.3$, standard ABC underestimates ϵ_{true} by about 0.2 and gives a diffuse final distribution, while FTLE-ABC gives a more narrow distribution centred on the true parameter value. These results suggest FTLE-ABC offers an even greater advantage over standard ABC in high chaoticity regimes when structures are spatially traversing the flow. This also demonstrates that FTLE-ABC can indeed deal with spatially traversing chaotic structures and may provide an advantage over standard methods for parameter estimation in geophysical flows with such structures. Due to this, only the model in the moving frame, where structures are spatially traversing, is considered for the comparison to the PF in Section 4.4.

4.4 Standard PF versus FTLE-ABC

Comparison of FTLE-ABC to a standard particle filter requires quite a different consideration of the parameters involved. The main weakness of a PF in the context of parameter estimation is associated with degeneracy. If a PF settles on an inaccurate parameter dis-

tribution initially and becomes degenerate it is unlikely that the algorithm will recover through further updates, due to constant re-sampling of the incorrect distribution. Hence, if the initial parameter distribution is quite broad *and* the PF is degenerate, the parameter estimate will be inaccurate. For Rossby wave flow in the moving frame initialised on a uniform grid the performance of the standard PF was found to be sensitive to noise in the initial positions of particles (σ_3) and the amount of time between observations (ΔT). When both are sufficiently large the PF becomes degenerate and re-samples on a significant proportion of the total number of assimilation steps. A large ΔT does not necessarily mean the value the PF estimates for ϵ_{true} is inaccurate, but it does tend to induce degeneracy, so that it becomes difficult to add information to the distribution of ϵ_{true} at each update. This, coupled with a reasonably high level of initial noise, causes the PF estimates for ϵ_{true} to become quite inaccurate.

In Section 4.3 it was shown that FTLE-ABC performs quite well for larger ΔT , suggesting this is a parameter regime for which FTLE-ABC may provide an advantage over a standard PF. FTLE-ABC should also be able to deal with a reasonable level of Gaussian noise in the initial conditions, as the stretching of fluid particles, as measured by the FTLE field, should not change significantly for zero-mean Gaussian noise. More generally, algorithms which can deal with initial uncertainty in tracer positions are advantageous in the context of geophysical parameter estimation, as tracer position data will naturally contain errors which may be magnified by chaotic behaviour. With these considerations in mind, varying ΔT and the initial noise added are the focus of the comparison of FTLE-ABC to the standard PF.

To assure the comparisons are done in a chaotic regime of Rossby wave flow, $\epsilon_{true} = 0.5$ is used for all the runs in this section. Also, as the results were quite mixed for FTLE-ABC when tracers were initialised only within gyres in Section 4.3, for the PF comparison tracers are only initialised on a uniform grid. All runs in this section use the low model dimension of $M = 4$, so that the PF does not require a computationally unreasonable number of ϵ realisations to avoid degeneracy (i.e. so that any degeneracy is induced by the magnitude of ΔT rather than the high dimension). Additionally, only the Rossby wave model in the moving frame is considered for the PF comparison, as it was established in the previous section that FTLE-ABC's performance is not effected by the presence of spatially traversing structures, and the moving frame scenario is more applicable to geophysical applications.

4.4.1 No initial noise and small versus large time steps

As a comparison to FTLE-ABC's distance versus ϵ plots, the PF weights w_j^i generated at each assimilation step are plotted against their corresponding parameter realisations

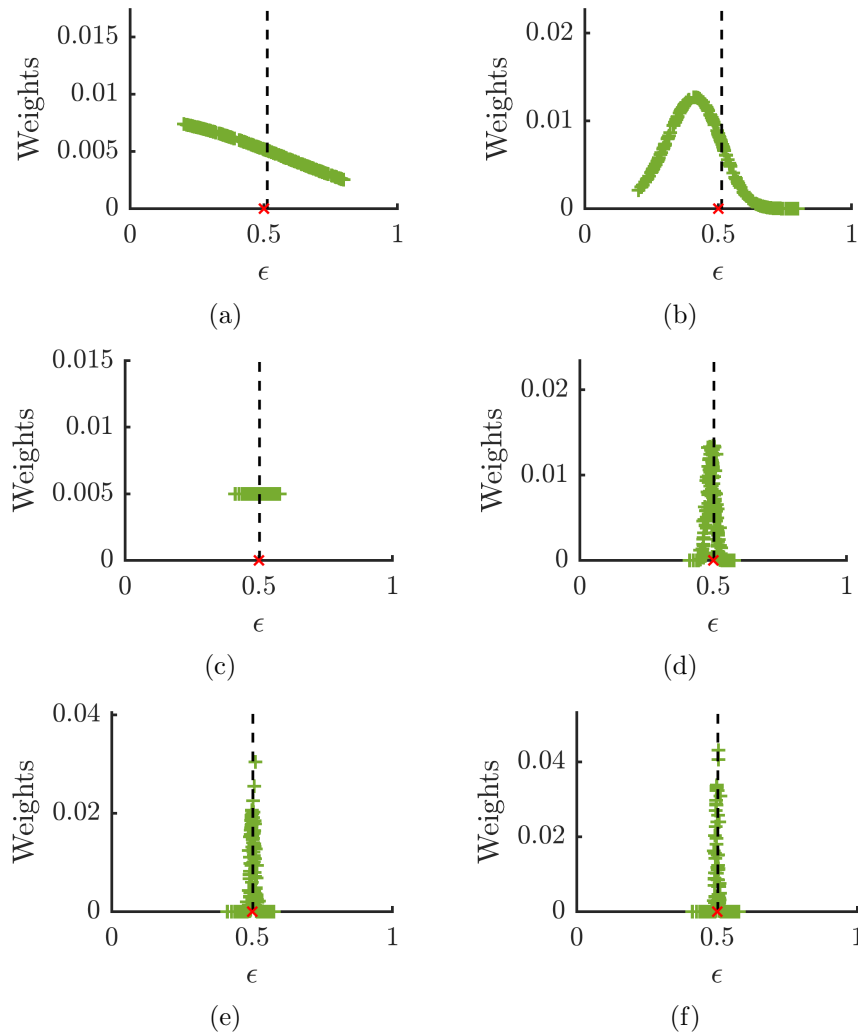


Figure 4.8: Comparison of 6 assimilation steps for the standard PF, in the moving frame with $T_f = 1$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.0$ and $N = 200$. Green crosses are particles, the black dashed line is the particle population average and the red cross is ϵ_{true} .

ϵ_j^i . The weights for the PF are equivalent to the distance measure for FTLE-ABC and ABC, as they provide a measure of the closeness of the PF's estimate for ϵ to the true value. For the PF, however, a maximum in the weights near ϵ_{true} , rather than a minimum, implies a good estimate, as higher weights imply a closer correspondence to ϵ_{true} . A weights distribution where all weights are equal regardless of ϵ implies that the PF has become degenerate and has resampled. In the PF plots in the following analysis, green crosses represent particles and black dashed lines represent the average ϵ of the particle population. In the FTLE-ABC plots, orange and blue diamonds again correspond to the

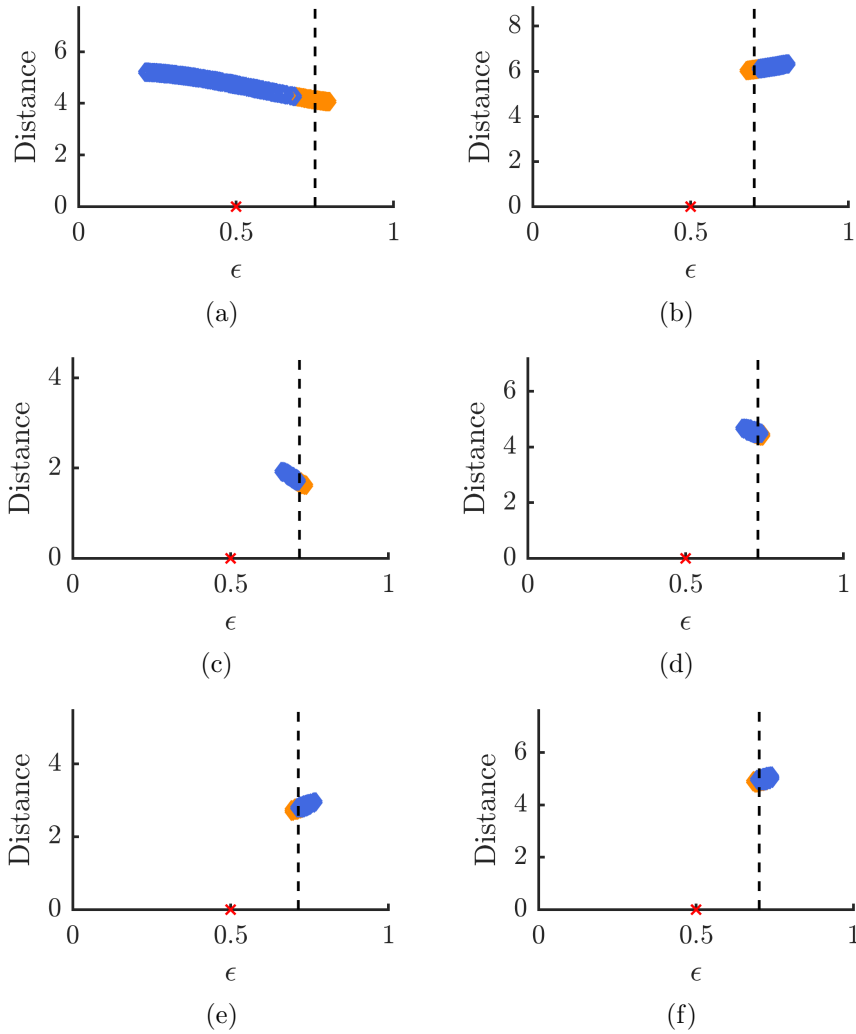


Figure 4.9: Comparison of 6 assimilation steps for FTLE-ABC, in the moving frame with $T_f = 1$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.0$ and $N = 200$. Blue and orange diamonds are particles, with orange representing the best 20% according to distance. The black line denotes the average of the best 20% and the red cross is ϵ_{true} .

particle population, with orange diamonds representing the best $\alpha = 20\%$ of particles according to distance, as in Section 4.3. The dashed black line represents the average ϵ for the best 20% of particles and ϵ_{true} is denoted with a red cross in all figures.

As a baseline comparison, conditions in which it is expected the PF will thrive are first considered, with no noise added initially to tracer positions. This scenario is compared for a long final time $T_f = 24$ and a short final time $T_f = 1$ for the same number of observations

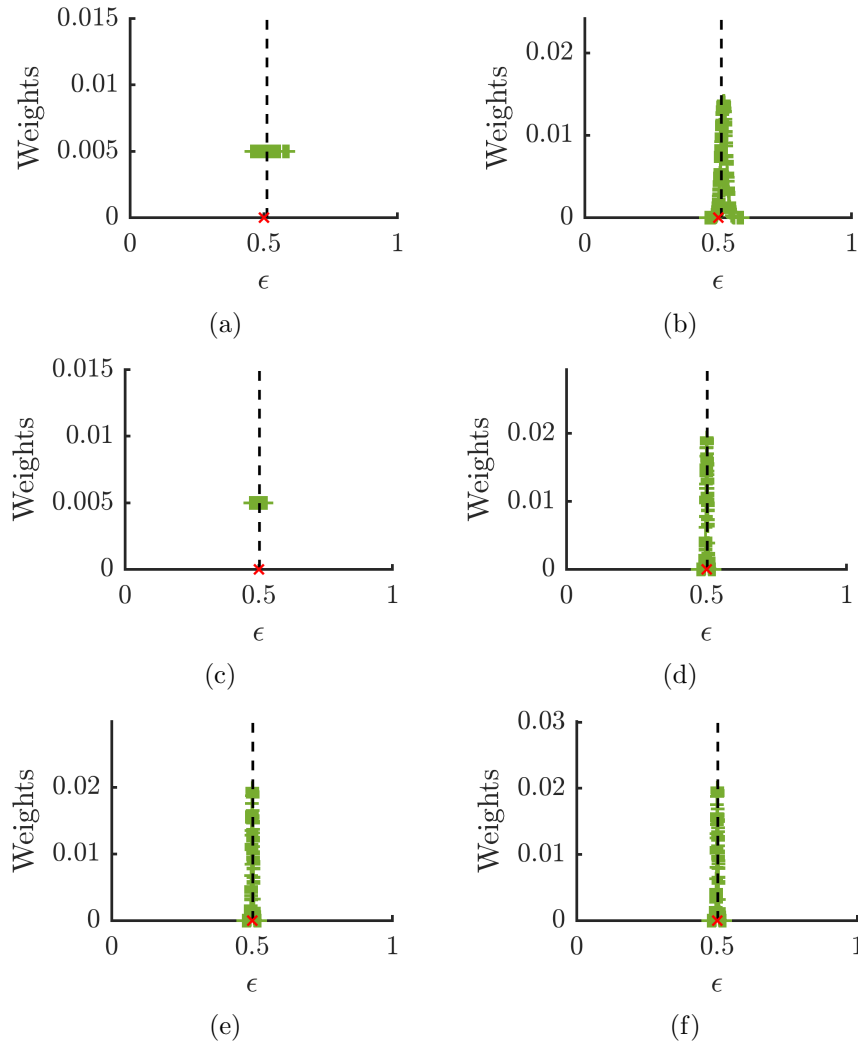


Figure 4.10: Comparison of 6 assimilation steps for the standard PF, in the moving frame with $T_f = 24$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.0$ and $N = 200$. Green crosses are particles, the black dashed line is the particle population average and the red cross is ϵ_{true} .

$J = 6$, so that in the former case the flow should be fairly chaotic in the final steps, while in the latter case the flow will not be particularly chaotic. With no initial noise added, however, the PF may still perform well for the chaotic case, as it may settle on the correct distribution for ϵ before the flow becomes too chaotic. This will provide a comparison for the following section, in which initialisation with added noise is considered. Figures 4.8 and 4.9 compare each step of the PF scheme and FTLE-ABC scheme with $\Delta T = 1/6$ while Figures 4.10 and 4.11 compare each step of the two schemes for $\Delta T = 4$. As can be seen from Figure 4.8, for the relatively small $\Delta T = 1/6$ the PF resamples at only one

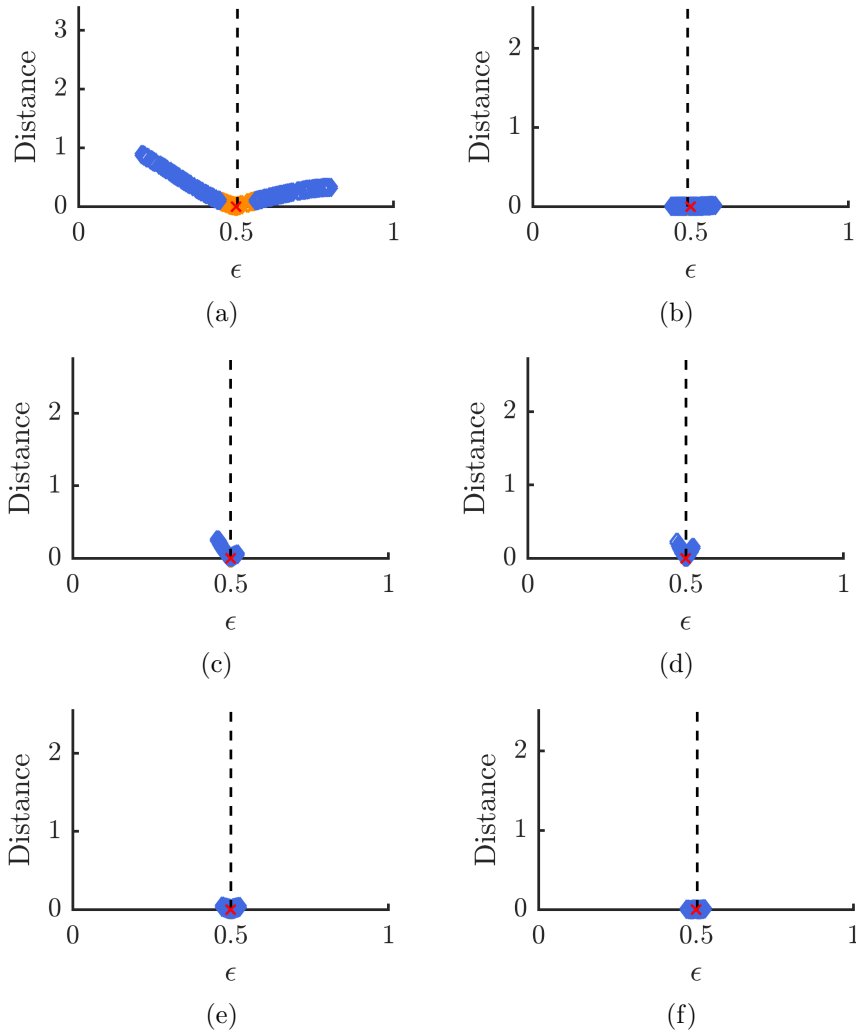


Figure 4.11: Comparison of 6 assimilation steps for FTLE-ABC, in the moving frame with $T_f = 24$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.0$ and $N = 200$. Blue and orange diamonds are particles, with orange representing the best 20% according to distance. The black line denotes the average of the best 20% and the red cross is ϵ_{true} .

assimilation step and forms a narrow maximum around $\epsilon_{true} = 0.5$. Comparatively, for these short time steps FTLE-ABC, shown in Figure 4.9, appears to fixate on an incorrect distribution for ϵ_{true} for all steps of the scheme. This is possibly due to the system being relatively stationary for small times, such that insufficient fluid-particle stretching occurs for the FTLE-ABC method to be informative, while the PF thrives in this minimal dynamics scenario. This explanation is indeed suggested by Figure 2.9 in Chapter 2, which shows that FTLE ridges will be poorly defined for this T_f . For the increased dynamics

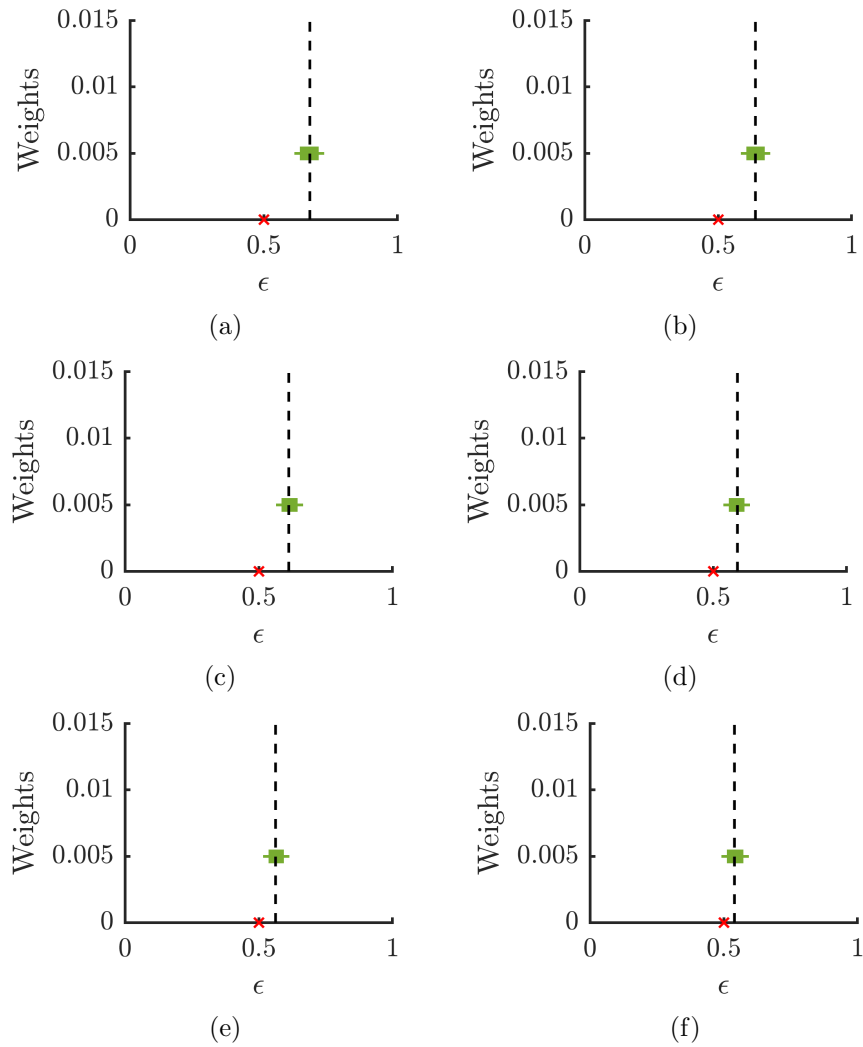


Figure 4.12: Comparison of 6 assimilation steps for the standard particle filter, in the co-moving frame with $T_f = 1$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.02$ and $N = 200$. Green crosses are particles, the black dashed line is the particle population average and the red cross is ϵ_{true} .

of the larger $\Delta T = 4$ scenario however, FTLE-ABC, shown in Figure 4.11, is able to quite accurately estimate ϵ_{true} and forms a narrow final distribution. The PF, shown in Figure 4.10, also estimates ϵ_{true} accurately for this larger final time with no initial noise, but must resample at more steps. Overall these results show that FTLE-ABC does not perform well where ΔT is too small to capture fluid particle stretching, but very well for larger ΔT , forming a narrow distribution about ϵ_{true} . The PF performs well in both regimes with no added initial noise, although the larger ΔT slightly increases the number

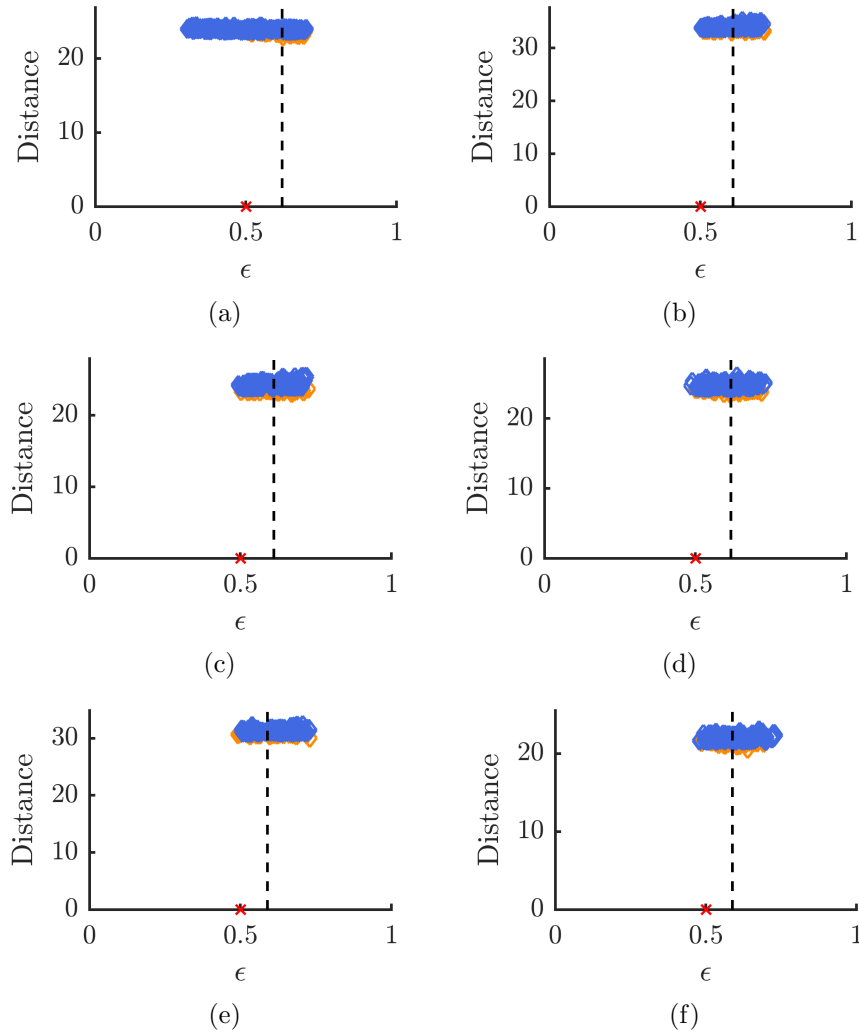


Figure 4.13: Comparison of 6 assimilation steps for FTLE-ABC, in the co-moving frame with $T_f = 1$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.02$ and $N = 200$. Blue and orange diamonds are particles, with orange representing the best 20% according to distance. The black line denotes the average of the best 20% and the red cross is ϵ_{true} .

of times it must resample.

4.4.2 Added initial noise and small versus large time steps

Here we consider the effect of adding zero-mean Gaussian noise in initial particle positions on the PF's performance compared to FTLE-ABC. It is expected that if the PF becomes

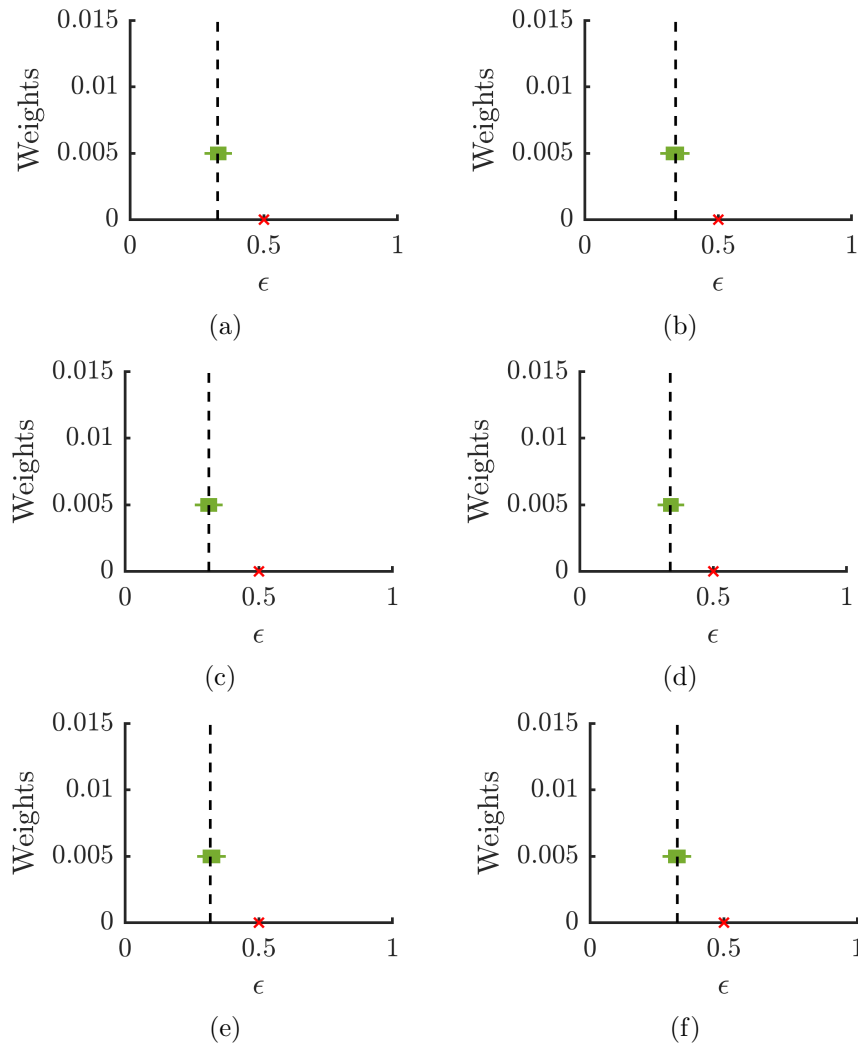


Figure 4.14: Comparison of 6 assimilation steps for the standard particle filter, in the co-moving frame with $T_f = 24$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.02$ and $N = 200$. Green crosses are particles, the black dashed line is the particle population average and the red cross is ϵ_{true} .

degenerate on an incorrect initial distribution it will not be able to recover, while information from the FTLE, on which FTLE-ABC relies, will not be severely effected. This is important as chaotic flows may magnify error in initial conditions as the model is run over time, such that a degenerate PF may drift increasingly far from an accurate estimate of a parameter, while a coherent structure method may recover from initial inaccuracy. Figures 4.12, 4.13, 4.14 and 4.15 again compare each step of the PF scheme and FTLE-ABC scheme for $\Delta T = 1/6$ and $\Delta T = 4$, but now with Gaussian noise of variance $\sigma_3 = 0.02$

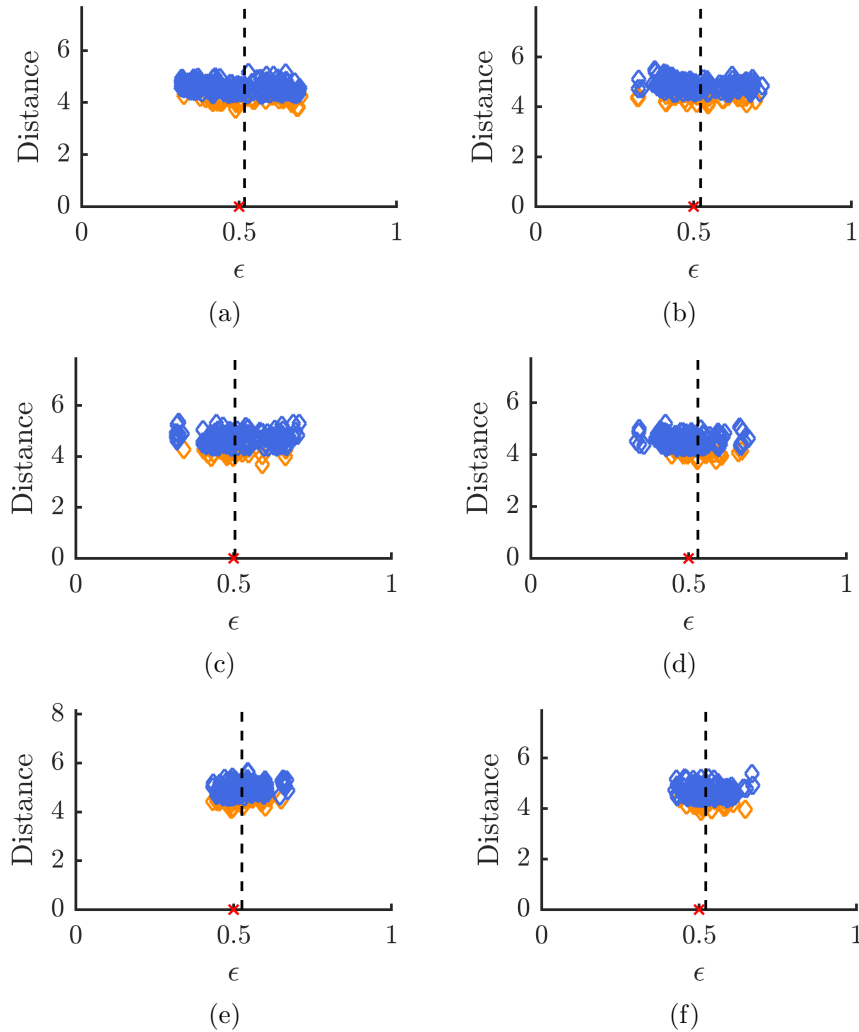


Figure 4.15: Comparison of 6 assimilation steps for FTLE-ABC, in the co-moving frame with $T_f = 24$, $J = 6$, $M = 4$, $\epsilon = 0.5$, $\sigma_3 = 0.02$ and $N = 200$. Blue and orange diamonds are particles, with orange representing the best 20% according to distance. The black line denotes the average of the best 20% and the red cross is ϵ_{true} .

added to initial tracer positions. When this small amount of Gaussian noise is added initially the performance of the standard PF changes considerably. For $\Delta T = 1/6$ in Figure 4.12 the PF now must resample in all of the assimilation steps, and only becomes close to ϵ_{true} in the final two steps. FTLE-ABC, shown in Figure 4.13, also does not estimate the correct ϵ_{true} with this small final time, tending to overestimate ϵ_{true} despite seeming to drift closer to the true value as the assimilation scheme progresses. As in the case where no noise is added initially, for this small final time the system is not sufficiently

dynamic for FTLE ridges to form and inform FTLE-ABC's estimate. For $\Delta T = 4$ the PF is again highly degenerate, shown in Figure 4.14, with resampling occurring on all assimilation steps. Due to the higher ΔT coupled with added noise, each assimilation step for the PF now shows estimates for ϵ_{true} around 0.1 lower than the true value in all steps, compared to 0.07 – 0.01 lower in the small ΔT with noise case. Comparatively, FTLE-ABC, shown in Figure 4.15, selects out the true parameter quite accurately at each assimilation step for $\Delta T = 4$, with the distribution becoming increasingly narrow about ϵ_{true} . As expected, FTLE-ABC is not severely effected by the added noise, and performs far better with added noise and a large final time than with added noise and a small final time. This large final time would allow FTLE ridges sufficient time to develop, so that FTLE-ABC may fully exploit them, while the increased dynamics of the larger final time worsens the inaccurate estimates of the PF when noise is added initially.

Overall, the investigations in Chapter 4 suggest that ϵ and ΔT have a particularly significant impact on the performance of FTLE-ABC compared to standard ABC, and that the algorithm shows the greatest advantage when in the moving frame of Rossby wave flow with the uniform initialisation scheme. For comparison to the standard PF, the amount of noise in initial tracer positions has a significant impact on the performance of the PF, but little impact on the performance of FTLE-ABC. These findings motivate Chapter 5, which uses a statistical measure of error, the root-mean squared error, to quantitatively compare the performance of FTLE-ABC to the benchmark algorithms in a range of ϵ and ΔT regimes.

Chapter 5

Quantifying FTLE-ABC's performance in chaotic regimes

Chapter 4 found that the quantities ϵ_{true} and ΔT have a significant impact on the performance of FTLE-ABC compared to the benchmark algorithms. To better quantify the relationship between ϵ_{true} , ΔT and FTLE-ABC's performance in parameter estimation, Chapter 5 introduces a statistical measure for the expected performance of each algorithm given many runs: the root-mean squared error (RMSE). The RMSE is used to quantitatively compare the performance of FTLE-ABC to that of the standard algorithms for various ϵ_{true} and ΔT regimes.

5.1 The root-mean squared error

The RMSE is a common measure of statistically relevant error arising from numerical processes. It is essentially a measure of the average distance between the approximate values of a parameter that a numerical process generates and the true parameter value. In the case of the algorithms employed here to estimate the true value of ϵ for Rossby wave flow, the RMSE considers the difference between ϵ_{true} , and K realisations of the DA scheme. From each realisation of the DA scheme, the quantity $\bar{\epsilon}(\epsilon_J)_k$ is calculated, where k is the realisation number ($k = 1, \dots, K$), and J is the index corresponding to the final time step of the assimilation scheme. The quantity $\bar{\epsilon}(\epsilon_J)_k$ is the average value of ϵ over the N particles generated by the k th independent run of the parameter estimation algorithm, calculated as

$$\bar{\epsilon}(\epsilon_J)_k = \frac{1}{N} \sum_{i=1}^N \epsilon_{i,J},$$

where i indexes members of the particle ensemble, J is the final assimilation time step and k indexes the number of times the parameter estimation scheme is run independently.

The form of the RMSE is then

$$\chi(\epsilon_{true}, \Delta T) = \sqrt{\frac{1}{K} \sum_{k=1}^K \left(\bar{\epsilon}(\epsilon_J)_k - \epsilon_{true} \right)^2}, \quad (5.1)$$

where ΔT is the assimilation time step.

Equation (5.1) gives a measure of how effectively a numerical scheme estimates ϵ_{true} over K independent runs of the scheme. This equation provides a performance comparison between the three algorithms as it measures the average error in an algorithm's estimate for ϵ_{true} over a large number of runs of the algorithm, rather than comparing single runs as in Chapter 4. To compare the RMSEs of the algorithms for a range of ΔT and ϵ_{true} , in Sections 5.2 and 5.4 RMSE contours against ΔT and ϵ_{true} are plotted for standard ABC, a standard PF and FTLE-ABC, denoted χ_{ABC} , χ_{PF} and χ_{F-ABC} respectively. In Sections 5.3 and 5.5 contour plots of the ratios χ_{F-ABC}/χ_{ABC} and χ_{F-ABC}/χ_{PF} are also plotted against ΔT and ϵ_{true} , denoted $\tau_{F/ABC}$ and $\tau_{F/PF}$ respectively. These ratio contours indicate regions where FTLE-ABC produces a lower RMSE than a standard method, as the ratio will be less than 1 in regions where this is true. The magnitude of the ratio also indicates the proportion by which FTLE-ABC provides a reduction in RMSE for regions where the ratio is less than 1.

Although the RMSE plots in the following sections generally support the findings of Chapter 4, they also reveal a complex relationship between FTLE-ABC's performance, ϵ , ΔT and the chaotic timescale, as estimated by the $1/FTLE_{med}$ curve from Section 2.8. As discussed in detail in the following sections, FTLE-ABC tends to give a lower RMSE than the standard methods for regions of ΔT and ϵ which lie below the $1/FTLE_{med}$ curve, or where the $1/FTLE_{med}$ curve asymptotes.

5.2 Comparing RMSEs: standard ABC comparison

For the comparison to standard ABC, RMSE contour plots were produced using $N = 200$ particles, a spatial dimension of 100, initially added Gaussian noise in tracer positions with variance $\sigma_3 = 0.01$ and a final time of $T_f = 24$ for the assimilation scheme. As it was demonstrated in Chapter 4 that the performance of FTLE-ABC is independent of the choice of frame, the moving frame of Rossby wave flow, for which structures may spatially traverse the flow domain, is used for all the runs considered in Chapter 5. Additionally, Chapter 4 suggests FTLE-ABC struggles with non-uniform initialisation, so the uniform initialisation scheme is used for all runs in Chapter 5.

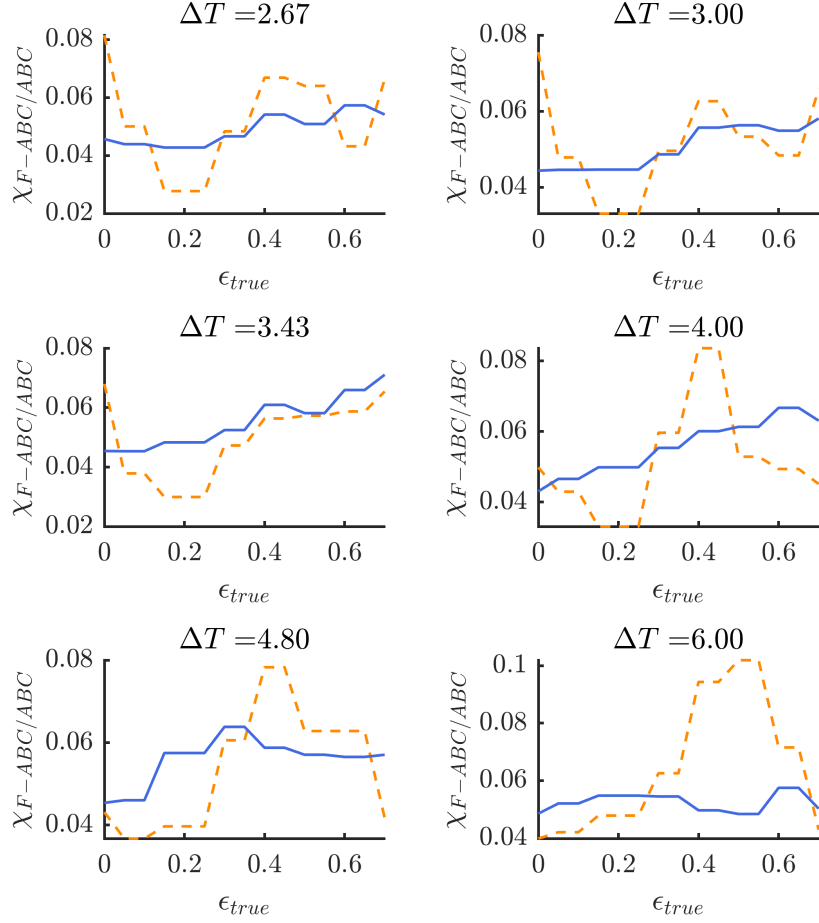


Figure 5.1: Plots of χ_{F-ABC} (orange dashed) and χ_{ABC} (blue filled) over a range of ϵ_{true} values for a selection of ΔT . A fairly consistent peak in RMSE for FTLE-ABC occurs around $\epsilon = 0.5$ and trough around $\epsilon = 0.2$.

To plot the grid of RMSE contours with respect to ΔT and ϵ_{true} , the algorithms were run 180 times for each combination of $(\Delta T, \epsilon_{true})$ for $\Delta T \in [1, 12]$ and $\epsilon_{true} \in [0, 0.7]$. These 180 runs were then used to calculate an RMSE for each parameter combination to form a grid of RMSEs in $(\Delta T, \epsilon_{true})$ -space. Before analysing the RMSE contour plots themselves, notable elements of these plots can be seen from Figure 5.1, which shows the RMSE for both methods against ϵ_{true} for a selection of individual ΔT values used in the contour plots. Here, the best performance according to RMSE for FTLE-ABC occurs around $\epsilon \in [0.2, 0.3]$ for $\Delta T \in [2.67, 6]$. This is interesting as $\epsilon \in [0.2, 0.3]$ is around the turning point of the $1/FTLE_{med}$ plot for $T_f = 24$, where the slope of the plot begins to decay less quickly with ϵ as it approaches a minimum. Also, $\Delta T \in [2.67, 6]$ are all smaller than the chaotic timescale ($\approx 8 - 10$) for $\epsilon = 0.2$ and so would give a sufficient resolu-

tion of the chaotic dynamics. Possibly the RMSE for FTLE-ABC is small here because this range of ΔT is smaller than the chaotic timescale for $\epsilon_{true} \in [0.2, 0.3]$, but not too small for FTLE ridges to be well-defined, thus allowing key structures to be captured by FTLE-ABC.

A peak in RMSE for FTLE-ABC around $\epsilon = 0.5$ also consistently occurs in these plots for $\Delta T \in [2.67, 6]$, which is around the chaotic timescale minimising ϵ for $T_f = 24$. This value of ϵ possibly gives a peak in RMSE because it is where the chaotic timescale is smallest on the $1/FTLE_{med}$ curve for $T_f = 24$. With this small chaotic timescale (≈ 5), $\Delta T \in [2.67, 6]$ would not be much smaller than the chaotic timescale and hence may not sufficiently resolve the chaotic dynamics. Indeed, the peak grows in magnitude with growing ΔT , being smallest for $\Delta T = 2.67$ and largest for $\Delta T = 6$ (which would be greater than the chaotic timescale). Additionally, it was suggested in the analysis of Figure 2.9, where FTLE fields are plotted for various ϵ and T_f , that for $\epsilon = 0.5$ and $T_f = 24$ the flow appears to be transitioning between two structural formations, causing higher noise in the FTLE field, which could impact the RMSE for FTLE-ABC.

After this peak in RMSE, however, the RMSE curve for FTLE-ABC begins to decay again, particularly for $\Delta T > 3.43$. This is possibly because the $1/FTLE_{med}$ curve begins to asymptote to around 5 and so the chaotic timescale is no longer changing strongly with ϵ . It is not obvious why a chaotic timescale which is approximately constant in ϵ would improve FTLE-ABC's performance, but it is possibly linked to how the algorithm relies on realisations of ϵ near ϵ_{true} to form an estimate of ϵ_{true} . If the chaotic timescale changes quickly with ϵ , a particular ΔT chosen for the assimilation scheme may not be small enough to capture the system dynamics for all of the ϵ realisations considered, which will cause varying error in the FTLE fields and the quality of FTLE ridges depending on the ϵ realisation. This may bias the scheme towards certain ϵ values and impact the RMSE. To investigate this, the impact of changing the width of the distribution from which ϵ realisations are drawn on the RMSE could be further examined. Additionally, it would be interesting to investigate whether the $1/FTLE_{med}$ curve plays a similar role in shaping the RMSE contours for FTLE-ABC for different final times T_f , which would indicate if the trends noted here extend more generally. For standard ABC the RMSE grows slowly but steadily with ϵ for $\Delta T \in [3, 4]$, but seems to asymptote to around 0.05 for other times. Overall these plots suggest a complex relationship between the RMSE, chaotic time scale, ΔT and ϵ for FTLE-ABC, which in particular combinations of these parameters give the algorithm a performance advantage over standard ABC.

Figure 5.2 shows RMSE contours for FTLE-ABC and standard ABC, created by plotting curves like those in Figure 5.1 as a contour plot in ϵ and ΔT . For FTLE-ABC, the RMSE plot peaks predominantly around $\Delta T = 6$ for $\epsilon_{true} \in [0.4, 0.55]$ which are values




Figure 5.2: RMSE contours (a) χ_{F-ABC} and (b) χ_{ABC} against the assimilation time step ΔT and ϵ_{true} . The model is in the moving frame with particles initialised on a standard grid, for a final time of $T_f = 24$. The $1/FTLE_{med}$ vs. ϵ curve for $T_f = 24$ is also plotted in red in (a).

nearby the chaotic timescale minimising ϵ for $T_f = 24$ on the $1/FTLE_{med}$ curve (shown in red). As suggested by the plots in Figure 5.1, FTLE-ABC performs worse in regions for which the $1/FTLE_{med}$ curve is approaching the minimum chaotic timescale and ΔT is larger than or of a similar magnitude to this chaotic timescale. Possibly the FTLE calculations require ΔT large enough for ridges to be defined but not so large as to poorly resolve the chaotic dynamics. Once the $1/FTLE_{med}$ curve has begun to asymptote following this minimum, however, the RMSE of FTLE-ABC improves, again suggesting that the method performs badly in regimes where the chaotic timescale is changing quickly with ϵ , but may recover if the chaotic timescale is fairly constant, even for fairly large ΔT . For standard ABC, an RMSE peak spreads over $\epsilon_{true} > 0.3$ for $\Delta T \in [3, 6]$. This peak also grows in magnitude fairly smoothly with ϵ_{true} , with the lowest values in RMSE tending to occur for the lowest ϵ values and higher RMSE values tending to occur for higher ϵ . This is expected, as in the lower ϵ regimes structural information is not as important and standard ABC may perform as well as FTLE-ABC, while such information should become significant for higher ϵ .

For FTLE-ABC the lowest RMSE region occurs for $\epsilon_{true} \in [0.1, 0.3]$ and $\Delta T \in [2, 6]$, which is entirely below the $1/FTLE_{med}$ curve. This again suggests that for lower ϵ , where the chaotic timescale is changing quickly with ϵ , FTLE-ABC performs better when ΔT is smaller than the chaotic timescale. FTLE-ABC's RMSE is also fairly low around $\epsilon_{true} \geq 0.6$ for most ΔT , a region where the $1/FTLE_{med}$ curve is asymptoting. Comparing the two contour plots, the behaviour of the RMSE for FTLE-ABC with ϵ_{true} and ΔT is much more varied than for standard ABC, and an overall trend is difficult to describe but seems linked to the chaotic timescale, as approximated by the $1/FTLE_{med}$ curve. Due to this complex behaviour, parameter regions for which FTLE-ABC gives a smaller RMSE than standard ABC are much more easily deduced from the ratio contour plots described in the following section.

5.3 RMSE ratios: standard ABC comparison

In developing a new LCS-based data assimilation method we are most interested in regimes in which the method out performs the standard methods *and* by how much the new method provides a performance advantage. In the context of this investigation we are interested in regimes in which FTLE-ABC provides a smaller RMSE than a standard method and how much smaller the method's RMSE is compared to a standard method. With this motivation, the contour plots of RMSE from the previous section were also plotted as ratio contours, with the contour colour indicating the magnitude of the ratio $\chi_{F-ABC}/\chi_{ABC} = \tau_{F/ABC}$. A ratio $\tau_{F/ABC} > 1$ implies a worse performance as measured by the RMSE than the standard ABC algorithm, while a ratio $\tau_{F/ABC} \leq 1$, implies a per-

formance equal to or better than standard ABC. To delineate these two scenarios visually, the contours are plotted so that a green gradient implies a ratio $\tau_{F/ABC} < 1$, a red gradient implies a ratio $\tau_{F/ABC} > 1$ and white space implies the borderline case $\tau_{F/ABC} = 1$. A colour bar is also shown on the plots to detail the magnitude of the ratio.

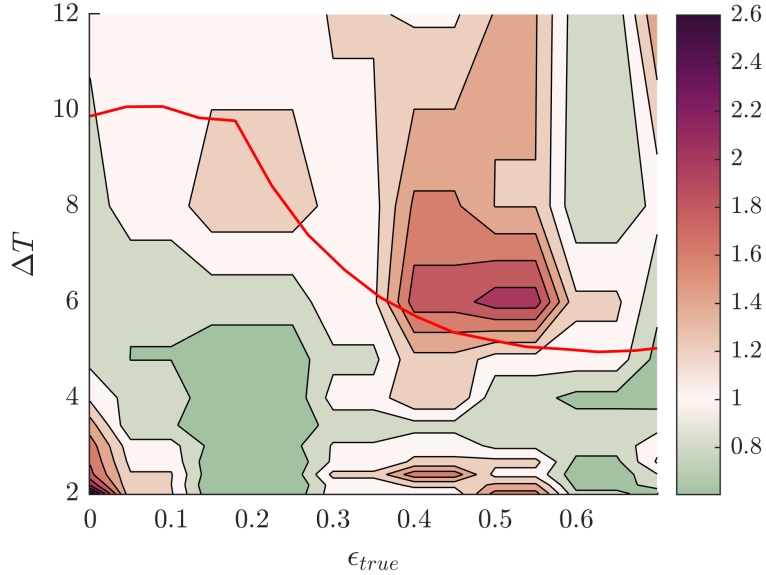


Figure 5.3: Contours of $\tau_{F/ABC}$ for the model in the moving frame, for a final time of $T_f = 24$, initialised on a uniform grid with $M = 100$ tracers. The red line is the $1/FTLE_{med}$ vs ϵ plot for $T_f = 24$.

The contour plot in Figure 5.3 shows the contours of the ratio $\tau_{F/ABC}$ with ϵ_{true} and ΔT , along with the $1/FTLE_{med}$ line for $T_f = 24$ overlaid in red. In this plot regions for which FTLE-ABC gives a lower RMSE than standard ABC become more obvious, as well as those for which FTLE-ABC gives a much higher RMSE. For low ϵ_{true} (0.1-0.3), and a low to medium range of ΔT (2-6), FTLE-ABC gives an RMSE 0.8 to 0.5 times that of standard ABC. This region is entirely for ΔT smaller than the approximate chaotic timescale (below the $1/FTLE_{med}$ line) and for ϵ_{true} not near the chaotic timescale minimising ϵ identified in Section 2.8 (≈ 0.5). This agrees well with the postulation made in Section 5.2 that for smaller ϵ , where the $1/FTLE_{med}$ curve has not yet asymptoted, FTLE-ABC performs well when ΔT is smaller than the chaotic timescale but still large enough for the formation of FTLE ridges.

There is a portion of the $\tau_{F/ABC}$ contour plot for which standard ABC and FTLE-ABC perform similarly or FTLE-ABC performs worse than standard ABC, which lies predominantly above the red $1/FTLE_{med}$ line and is focused around the chaotic timescale

minimising ϵ for $T_f = 24$. There is also a region around $\epsilon_{true} \geq 0.6$ for which FTLE-ABC gives an RMSE 0.5 to 0.8 times that of standard ABC for almost all ΔT , which occurs where the $1/FTLE_{med}$ line begins to asymptote and no longer change significantly with ϵ . Although not done here, it would be interesting to extend the range of ϵ_{true} to see if this region where FTLE-ABC achieves a lower RMSE for most ΔT continues as the $1/FTLE_{med}$ line asymptotes for $\epsilon > 0.7$. Overall, the ratio contours suggest that for lower ϵ , where the chaotic timescale is still changing rapidly with ϵ , FTLE-ABC requires a ΔT smaller than the chaotic timescale to out perform standard ABC but, when the chaotic timescale begins to become constant with ϵ , FTLE-ABC seems to out perform standard ABC for a much wider range of ΔT . As discussed in section 5.2, this may be linked to the width of the interval from which the ϵ realisations are drawn in FTLE-ABC, although further investigations would be required to confirm this.

5.4 Comparing RMSEs: standard PF comparison

For the standard PF comparison again a final time of $T_f = 24$ was used for both schemes, but with a spatial dimension of 36 (due to the PFs tendency to be degenerate for high model dimensions). In these runs $N = 50$ particles were used and Gaussian noise in initial drifter positions with variance $\sigma_3 = 0.2$. In addition to posing challenges for the PF, this high initial error regime mimics aspects of real-world scenarios. A common issue in modeling drifters is that uncertainty in drifter positions often grows large as the model is simulated, worsened in the presence of numerical error, model error and chaotic behaviour. Rather than adding, for example, additional model error at each time step, we mimic the effect of this by using large initial uncertainty in our chaotic model. Additionally, the lower number of particles used ($N = 50$), which may induce degeneracy in the PF, is more practical given the high-dimensionality of many geophysical models. Using these conditions for comparison to the PF thus allows insight into FTLE-ABC's performance when the number of realisations used, N , is reduced (giving a more computationally efficient algorithm) and when the error in drifter positions grows large as the model is simulated.

As in Section 5.2, to plot the RMSE contours for the PF comparison the algorithms were run 180 times for each combination of $(\Delta T, \epsilon_{true})$ for $\Delta T \in [1, 12]$ and $\epsilon_{true} \in [0, 0.7]$. Before analysing the resulting contours, notable elements of the RMSE trends for the PF and FTLE-ABC can be discerned from Figure 5.4, which shows plots of the RMSE against ϵ for a selection of ΔT values. Here, for $\Delta T \in [2, 8]$ and $\epsilon_{true} > 0.4$, FTLE-ABC tends to give a lower RMSE than the standard PF, and its RMSE tends to be minimised around $\epsilon = 0.5$. This is quite different from the comparison to standard-ABC, in which FTLE-ABC's RMSE tended to peak around $\epsilon = 0.5$. FTLE-ABC's RMSE now peaks around

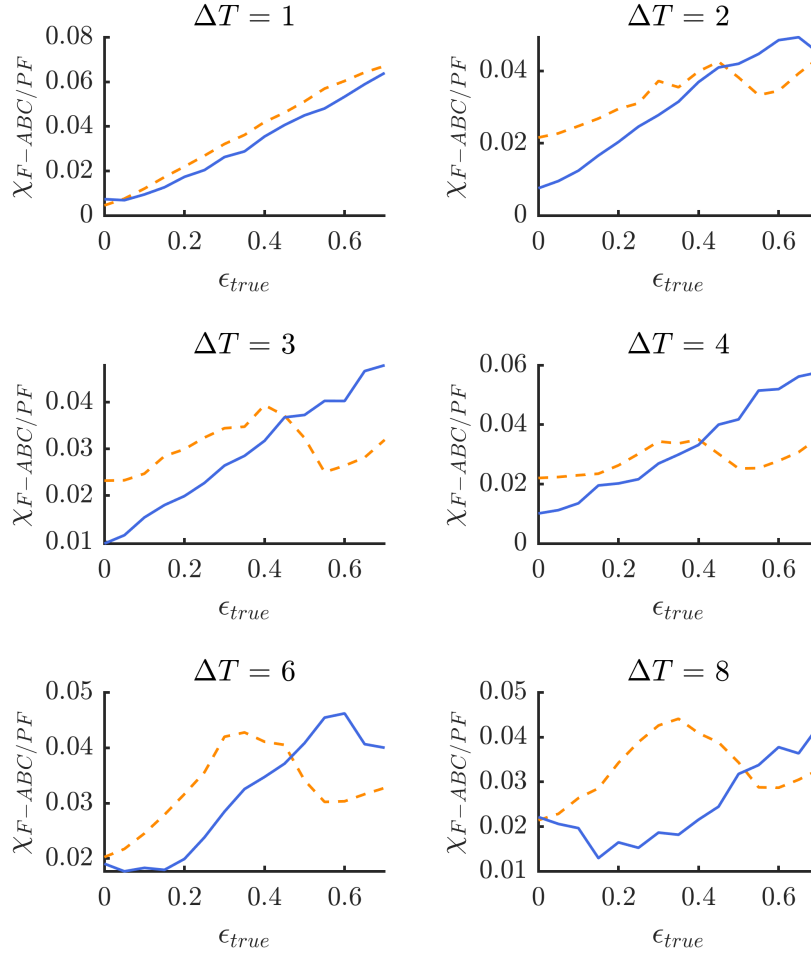


Figure 5.4: Plots of χ_{F-ABC} (orange dashed) and χ_{PF} (blue filled) over a range of ϵ_{true} values. As can be seen, for $\Delta T \in [2, 8]$ and $\epsilon_{true} > 0.4$, FTLE-ABC tends to out-perform the standard PF, while for lower ΔT and ϵ the PF performs better.

$\epsilon = 0.3$ for $\Delta T \in [3, 8]$, possibly due to the high level of noise added initially causing conditions similar to higher ϵ regimes in the lower initial noise case. Despite this difference, the ϵ value at which the FTLE-ABC and PF RMSE curves tend to cross over, with FTLE-ABC's RMSE becoming lower than that of the PF, is around $\epsilon \in [0.4, 0.5]$, which is around the ϵ value for which the $1/FTLE_{med}$ curve begins to flatten for $T_f = 24$. This is similar to the FTLE-ABC comparison, in which regions where the chaotic timescale is not changing rapidly with ϵ give FTLE-ABC a lower RMSE for a wider range of ΔT . Unlike in the ABC comparison however, FTLE-ABC's RMSE appears to grow again following this minimum, although for $\Delta T \in [3, 8]$ the RMSE grows more slowly than that for the standard PF and so stays lower than its RMSE.




Figure 5.5: Contours of RMSE for (a) χ_{F-ABC} and (b) χ_{PF} against the assimilation time step ΔT and ϵ_{true} , for the model in the moving frame with particles initialised on a standard grid, for a final time of $T_f = 24$.

RMSE contours against ϵ and ΔT are shown in Figure 5.5 for FTLE-ABC (5.5a), and

the standard PF (5.5b). As suggested by the individual RMSE curves, the ϵ_{true} region for which FTLE-ABC performs the worst occurs around $\epsilon_{true} \in [0.3, 0.4]$ for $\Delta T \in [6, 12]$, compared to around $\epsilon_{true} = 0.55$ and $\Delta T \in [4, 8]$ for the comparison to standard ABC. The method again performs fairly well for $\epsilon_{true} \in [0, 0.3]$ for most ΔT , but also performs well for a region of higher $\epsilon \in [0.4, 0.65]$ for $\Delta T \in [2, 8]$. This trough in the RMSE is notable as it occurs in the region of ϵ_{true} and ΔT for which the $1/FTLE_{med}$ curve begins to asymptote for $T_f = 24$, as suggested by the plots in Figure 5.4. Indeed, the RMSE is generally smaller for FTLE-ABC below regions where the $1/FTLE_{med}$ curve is flat.

Comparing to Figure 5.5b, the trend in RMSE with ΔT and ϵ_{true} is very different for the standard PF, which predominantly shows changes in RMSE with the magnitude of ϵ_{true} . The PF gives a lower RMSE when ϵ_{true} is small and higher RMSE when ϵ_{true} is large, which is expected as the higher ϵ regimes will be more chaotic and induce degeneracy in the PF, which will have initially settled on an inaccurate distribution for ϵ . Interestingly, the RMSE magnitude does not change significantly with ΔT for the PF, although for higher ϵ , higher ΔT produces a slightly lower RMSE than lower values. This seems to contradict the findings of Chapter 4, but is perhaps because larger time steps with higher chaoticity allow the PF's distribution for ϵ to become wider and hence more able to move away from an incorrect distribution it has settled on initially due to being degenerate. For lower ϵ , however, and $\Delta T \leq 6$ the RMSE does increase slightly with ΔT , which agrees with the parameter regimes of Chapter 4. Overall these trends support the postulation in Chapter 4 that for approximately stationary scenarios, which would occur for small ϵ , the PF should perform similarly or better than FTLE-ABC, while for larger ϵ we would expect FTLE-ABC to provide an advantage.

5.5 RMSE ratios: standard PF comparison

As in Section 5.3, to discern regions of ΔT and ϵ_{true} for which FTLE-ABC's RMSE is better than that of the standard PF, and the magnitude of the performance gain, contours of the ratio $\chi_{F-ABC}/\chi_{PF} = \tau_{F/PF}$ are also plotted. Figure 5.6 shows the contours of $\tau_{F/PF}$ with the $1/FTLE_{med}$ curve overlaid in red, and the same colour gradient to indicate the magnitude of the ratio as in Section 5.3. Although quite different from the ratio contours generated for the standard ABC comparison, there are subtle similarities to the $\tau_{F/ABC}$ contours in Section 5.3. There is a region around $\epsilon \in [0.5, 0.7]$ for which FTLE-ABC's RMSE is around 0.8-0.5 times that of the standard PF for most $\Delta T < 12$. This is similar to the band for $\epsilon > 0.5$ in the ABC comparison but slightly wider, which occurs for regions where the $1/FTLE_{med}$ curve asymptotes. The region around $\epsilon \in [0.1, 0.3]$ and small ΔT where FTLE-ABC's RMSE was smallest in the ABC comparison, however, is not present here. Instead there is a small band for $\Delta T \in [8, 12]$ and $\epsilon \in [0, 0.1]$ for which




Figure 5.6: Contours of $\tau_{F/PF}$ for the model in the moving frame, for a final time of $T_f = 24$, initialised on a standard grid for $M = 36$ tracers. The red line is the $1/FTLE_{med}$ vs ϵ plot for $T_f = 24$.

FTLE-ABC's RMSE is smaller than that of the PF. This is possibly due to the higher magnitude of initial noise causing the FTLE to require a larger ΔT to produce well-defined ridges for the smaller ϵ regimes. Interestingly, the two parameter regions for which the overlaid $1/FTLE_{med}$ curve flattens approximately line up with the regions of the $\tau_{F/PF}$ plot where FTLE-ABC gives a lower RMSE than the PF. Again this is possibly due to the width of the distributions from which ϵ realisations are drawn in the FTLE-ABC algorithm and how quickly the chaotic timescale is changing with ϵ within the range of these distributions. Whether this observation is a general trend could be investigated by comparing the RMSE contours to $1/FTLE_{med}$ curves for different T_f values and different widths of the distributions used for ϵ realisations, as mentioned in Section 5.3.

Overall, although Chapter 5 confirms some of the trends suggested by the initial investigations in Chapter 4, a much more complex relationship between the algorithm's RMSE, ΔT and ϵ is suggested by the findings of Sections 5.3-5.5, with the $1/FTLE_{med}$ curve with ϵ potentially giving some insight into this relationship. For the comparison to standard ABC, there are distinct parameter regions for which the method outperforms standard ABC, which depend strongly on both the value of ϵ_{true} and ΔT and seem to occur predominantly below the $1/FTLE_{med}$ curve or where it is asymptoting. The method particularly struggles for $\epsilon_{true} \in [0.4, 0.55]$, which corresponds to the chaotic timescale minimising ϵ_{true} from the $1/FTLE_{med}$ curve for $T_f = 24$. FTLE-ABC may also outper-

form the standard PF for particular regions of ΔT and ϵ , which seem to correspond to where the $1/FTLE_{med}$ curve flattens. Although the $1/FTLE_{med}$ curve has some predictive power when it comes to the error structure of FTLE-ABC, the curves tend to suggest where FTLE-ABC's RMSE will be better but not necessarily where FTLE-ABC will give a lower RMSE than a standard method (i.e. the χ_{F-ABC} contours tend to line up better with the $1/FTLE_{med}$ curves than the τ_{F-ABC} contours). Thus it seems that the chaoticity (as estimated by $1/FTLE_{med}$) seems necessary but not sufficient for better performance by FTLE-ABC. These findings, their interpretations and potential further research directions are discussed in Chapter 6.

Chapter 6

Summary, Discussion and Conclusions

Chapter 6 summarises the results of this thesis and discusses interpretations, limitations and potential future research avenues. The technique developed, FTLE-ABC, is found to outperform the benchmark algorithms for particular regions of ϵ_{true} and ΔT . These regions are possibly related to the rate of change of the chaotic timescale with ϵ_{true} , as measured by the $1/FTLE_{med}$ curve, although further investigations would be needed to confirm this. Overall the results presented here add a frame-independent, non-linear method to a growing literature in techniques amalgamating LCS and LaDA methods to improve predictions in chaotic geophysical systems.

6.1 Summary: LaCoDA and FTLE-ABC

The primary aim of this thesis is to build on current research in Lagrangian coherent data assimilation methods for parameter estimation. In particular, this thesis develops a Lagrangian coherent structure-based data assimilation technique aimed at overcoming issues associated with parameter estimation in highly chaotic geophysical systems. The particular LCS method used here is the Finite Time Lyapunov Exponent, a widely used, versatile and frame-independent technique commonly implemented to track structures in Lagrangian oceanographic data. The FTLE is used within an ABC-type data assimilation scheme for parameter estimation, denoted FTLE-ABC, which is implemented to estimate the ϵ parameter of Rossby wave flow for various parameter regimes.

The development, theory and utility of Lagrangian coherent structure methods in geophysical systems is discussed in Chapter 2, with an emphasis on how such methods seek to extract transient yet coherent structures extant in chaotic, time-dependent fluid flows. Data assimilation methods, theory and applications are detailed in Chapter 3, with an em-

phasis on Lagrangian data assimilation and the importance of such techniques in wholly utilising the rich information available in Lagrangian data. This leads to Chapter 4, which develops a formal algorithm for FTLE-ABC and compares the algorithm's performance to the benchmarks of a standard particle filter and a standard ABC scheme. Here it is found that ΔT , the assimilation time step, and ϵ_{true} , the true value of the parameter being estimated, appear to most strongly impact FTLE-ABC's parameter estimation performance. Chapter 5 details further investigations into how ϵ and ΔT influence the new algorithm's performance, which is quantified by the RMSE. The predominant findings of Chapters 4 and 5 are summarised in Sections 6.1.1 and 6.1.2 below.

6.1.1 FTLE-ABC versus standard ABC

In Chapter 4 individual runs of the FTLE-ABC scheme and standard ABC scheme were compared for the following different scenarios:

- For the low, middle and high values of ϵ : 0.3, 0.9 and 1.3. 'Low' and 'high' are measured here by the plots in Figure 2.8 for $T_f = 12$, which suggest 0.3 will have a longer chaotic timescale and 0.9 and 1.3 will have a very low chaotic timescale.
- For large $\Delta T = 4$, and smaller $\Delta T = 1.33$. This is again dictated by the plots in Figure 2.8 for $T_f = 12$, which suggest $\Delta T = 4$ is larger than the chaotic timescale for $\epsilon = 0.9$ and 1.3, while $\Delta T = 1.33$ is smaller than the chaotic timescale for $\epsilon = 0.3$, 0.9 and 1.3.
- For Rossby Wave flow in the co-moving frame and in the moving frame, for which the gyres are stationary and spatially traverse the flow domain respectively.
- For a uniform grid initialisation scheme, which evenly samples the flow domain, and for a non-uniform initialisation solely in the flow gyres.

For these individual runs it was found that FTLE-ABC generally provided a better estimate of the higher values of ϵ_{true} (0.9 and 1.3) when the larger time step of $\Delta T = 4$ was used. A potential cause of this behaviour is suggested by the plot for $T_f = 12$ in Figure 2.8, which implies a chaotic timescale around 6 for $\epsilon < 0.5$, less than 4 for $\epsilon = 0.9$ and around 3 for $\epsilon = 1.3$. Hence the smaller time step $\Delta T = 1.33$ is less than the estimated chaotic timescale for all the ϵ considered and should allow a sufficient resolution of the system dynamics. Comparatively, $\Delta T = 4$ is larger than the chaotic timescale for $\epsilon > 0.5$, suggesting that coherent structure information may be useful here. This is indeed the case for $\Delta T = 4$ and $\epsilon = 0.9$ and 1.3, with FTLE-ABC and standard ABC seeming to perform similarly for $\epsilon = 0.3$, but FTLE-ABC significantly out-performing standard ABC for the higher values of ϵ .

To capitalise on the additional structural information granted by the FTLE, however, it seems from the results in Chapter 4 that FTLE-ABC requires fairly uniform sampling of the flow domain. For scenarios in which FTLE-ABC out-performed standard ABC with the uniform grid initialisation, all such performance gains were lost when the non-uniform gyre initialisation scheme was used. This is perhaps due to the fact that the FTLE structural information relies on a contrast in FTLE magnitude between ridges and other portions of the field. With a non-uniform initialisation, higher magnitude ridges may have been biased towards certain portions of the domain, so that structural information was lost.

Despite this required specificity of initialisation scheme, the results in Chapter 4 show that FTLE-ABC performs almost identically when the Rossby wave model is in the moving and co-moving frames. This implies that the method is indeed able to successfully deal with spatially traversing structures, while standard ABC struggles in this scenario. Such a property is a useful addition to available structure-based parameter estimation methods, as methods such as the Principle Component Analysis technique used in Maclean et al. (2017) cannot deal with such structures.

In Chapter 5 potential performance advantages of FTLE-ABC over a standard ABC scheme are quantified by considering the RMSE for both parameter estimation methods over a range of ΔT and ϵ . The RMSEs are compared for the Rossby wave model in the moving frame with the uniform initialisation scheme for the following different scenarios:

- For $\epsilon_{true} \in [0, 0.7]$.
- For $\Delta T \in [1, 12]$ for $T_f = 24$.
- With a small amount of initial noise added ($\sigma_3 = 0.01$), and reasonably high model dimension $M = 100$.

From this more quantitative approach it was found that the relationship between the RMSE, ΔT and ϵ is complex, and may possibly be related to the rate of change of the chaotic timescale with ϵ , as estimated here by the $1/FTLE_{med}$ curve. RMSE contour plots for FTLE-ABC and standard ABC against ΔT and ϵ showed that FTLE-ABC tends to give a higher RMSE than standard ABC around $\epsilon_{true} = 0.5$ and $\Delta T \geq 6$, a region above where the $1/FTLE_{med}$ curve is decreasing strongly with ΔT . Additionally, $\epsilon_{true} = 0.5$ is around the chaotic timescale minimising ϵ identified for $T_f = 24$ in Figure 2.8, which may have caused this high RMSE for the algorithm. FTLE-ABC tends to give a lower RMSE than standard ABC around $\epsilon_{true} \in [0.1, 0.3]$ and $\Delta T \in [2, 6]$, a region entirely below the $1/FTLE_{med}$ curve. The algorithm also gives a lower RMSE than standard ABC for $\epsilon \geq 0.6$ for all ΔT , a region where the $1/FTLE_{med}$ curve is flat and appears to asymptote. As suggested in Chapter 5, a potential explanation for this is that, because ϵ

realisations are drawn from a range of values near ϵ_{true} , if the chaotic timescale is changing rapidly with ϵ the particular ΔT chosen for the scheme may not sufficiently resolve the dynamics for all of the ϵ realisations considered, giving varied error in the FTLE fields used to update the distribution. Such issues would not be present if the chaotic timescale is not changing rapidly with ϵ . A further investigation could be to examine how far this region expands for $\epsilon > 0.7$ as the $1/FTLE_{med}$ curve continues to asymptote, and also how changing the width of the distribution from which ϵ realisations are drawn impacts this trend.

More simplistically, these results agree well with the FTLE fields in Figure 6.1 (also shown in Figure 2.9 in Chapter 2), which show that the most well-defined ridges tend to occur for ranges of ϵ away from the chaotic timescale minimising ϵ . The plots in Figure 6.1 also suggest ridges are well-defined for ΔT large enough for structures to form, but not so large that the field becomes noisy. For example, for $\Delta T = 2$, ridges are most defined for $\epsilon \in [0.3, 0.9]$, while for the higher $\Delta T = 12$, structures are difficult to see when $\epsilon = 0.5$, but discernible for $\epsilon \in [0, 0.3]$. Additional structures then seem to reemerge for $\epsilon = 0.9$ surrounded by a background of noise, which would be a region of ΔT and ϵ for which the $1/FTLE_{med}$ curve asymptotes. How well-defined FTLE ridges are thus may be linked to the chaotic time scale and in turn could impact FTLE-ABC's performance. With this in mind, further insight into the regimes in which the algorithm's performance exceeds the standard methods could be gained by further investigating the impact of the $1/FTLE_{med}$ curve on the formation of FTLE ridges.

6.1.2 FTLE-ABC versus standard PF

Individual runs of the FTLE-ABC scheme and standard PF scheme were also compared for Rossby wave flow in the moving frame with the uniform initialisation scheme for the following different scenarios:

- For small final time $T_f = 1$, so that for $J = 6$, $\Delta T = 1/6$, and a larger final time $T_f = 24$, so that with $J = 6$, $\Delta T = 4$. These timescales are again smaller and larger than the approximate chaotic timescale respectively, as measured by the $1/FTLE_{med}$ curves in Figure 2.8.
- For no initial noise $\sigma_3 = 0$ and for higher initial noise $\sigma_3 = 0.002$.

For these individual runs it was found that for the very small final time of $T_f = 1$, giving $\Delta T = 1/6$, the PF outperforms FTLE-ABC, both when no initial noise is added and with small initial noise. This is possibly because the system is effectively stationary for such a short run time and the FTLE field ridges may be poorly defined for $T_f < 2$. More insight on this is available in Figure 6.1, where the FTLE fields are shown for different values

of ϵ and T_f . From this grid of FTLE fields it can be seen that for $\Delta T = 0.5$ the FTLE ridges are poorly defined for all ϵ , such that the FTLE-ABC algorithm would struggle to gain information from the FTLE. For the much larger final time of $T_f = 24$, giving $\Delta T = 4$, the PF becomes quite degenerate, although FTLE-ABC and the PF perform similarly in terms of selecting out the true parameter if no initial noise is added. With some initial noise and a large final time, FTLE-ABC generally outperforms the PF, most likely because the PF tends to become degenerate on an incorrect range of ϵ , while the FTLE structures which inform FTLE-ABC are not significantly effected by the added initial noise.

In Chapter 5 potential performance advantages of FTLE-ABC over a standard PF scheme are quantified by considering the RMSE of both parameter estimation methods for a range of ΔT and ϵ . The RMSE's for the model in the moving frame with the uniform initialisation scheme are compared for the following different scenarios:

- For $\epsilon_{true} \in [0, 0.7]$.
- For $\Delta T \in [1, 12]$ for $T_f = 24$.
- With a high amount of initial noise added $\sigma_3 = 0.2$ and lower model dimension $M = 36$.

Here, FTLE-ABC tends to give a lower RMSE than the standard PF in regimes of ϵ and ΔT for which the $1/FTLE_{med}$ curve flattens, and a higher RMSE than the standard PF for regions where the $1/FTLE_{med}$ curve has a high gradient. This somewhat agrees with the findings for the comparison to standard ABC, although the contour plots for the two comparisons have significant differences, most likely due to different model dimensions and magnitudes of noise added to tracer positions initially. The very high RMSE seen for FTLE-ABC in the standard ABC comparison case around $\epsilon_{true} = 0.5$ is now shifted to $\epsilon = 0.3$, possibly due to the high noise in initial positions causing conditions similar to a higher ϵ regime.

6.2 Further research pathways

The findings in this thesis suggest a complex relationship between the chaotic timescale and the performance of FTLE-ABC, with particular ΔT and ϵ regimes allowing the algorithm's performance to be better than the standard algorithms. To further elucidate these favourable parameter regimes an even finer grid in ΔT and ϵ for the RMSE plots generated in Chapter 5 could be considered, with a larger range of ΔT and ϵ values. A finer grid may resolve the fluctuations of the contour plots with ϵ_{true} , potentially giving

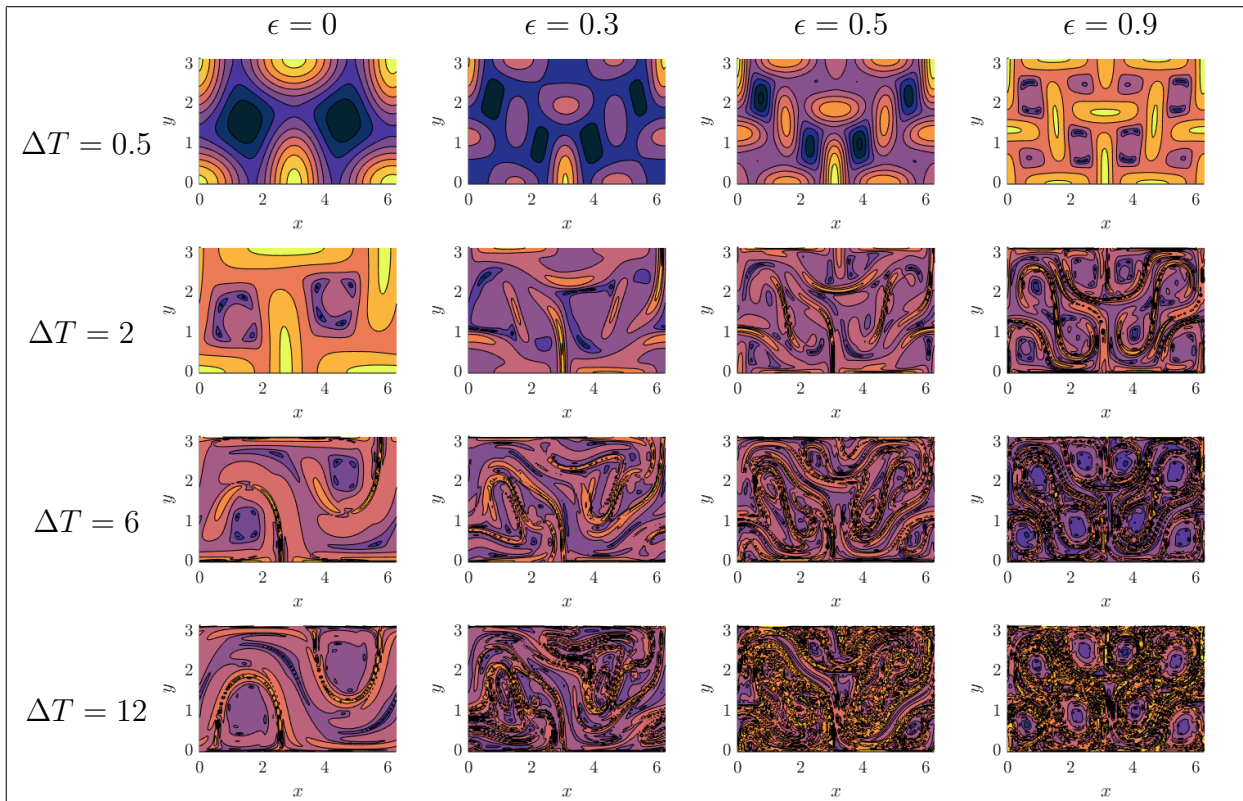


Figure 6.1: Table showing how the FTLE field varies with the calculation time $\Delta T = T_f - 0$ and the parameter ϵ .

deeper insight into the relationship between the RMSE and the $1/FTLE_{med}$ curves in Chapter 2. Further to this, computing RMSE contour plots against ΔT vs. ϵ for different T_f values and attempting to find a trend in the RMSE contour plots and their corresponding $1/FTLE_{med}$ curves with changing T_f may be a route to a deeper understanding of how the chaotic time scale effects the algorithm's performance. Additionally, although the chaoticity, as estimated by the $1/FTLE_{med}$ curve, was shown to have predictive power in suggesting parameter regions for which FTLE-ABC would give a lower RMSE, it could not always predict where FTLE-ABC would do better than a standard method. A deeper investigation into how well defined FTLE ridges are for particular ΔT and ϵ values could give further insight into FTLE-ABC's performance against the standard methods.

Another avenue for future research is the question of how to best target observations at each assimilation step. As shown by Salman et al. (2008), targeting observations by placing tracers to minimise uncertainty in observations can produce major benefits for LaDA schemes generally. A question which arises then is how to target observations to

optimise FTLE-ABC's performance. One way of doing this for FTLE-ABC might be to find methods of placing tracers such that the FTLE ridges are always well-defined. There are many possible ways to approach this problem, but it would require quite a different set of research questions to those considered here.

Another question not addressed here is how the new algorithm would perform on a real-world data set. Trialing the algorithm with the Rossby wave model but with experimental fluid data rather than simulated data would test how the algorithm deals with potentially non-Gaussian noise in observational data. As FTLE-ABC is based on ABC, which can deal with non-Gaussian likelihoods, and the FTLE, which is a non-linear method, the algorithm should theoretically be able to deal with such a scenario. This would be a considerable advantage over standard methods as many parameter estimation schemes are based on the Kalman Filter, which assumes Gaussian distributed noise. Additionally, trialing FTLE-ABC for parameter estimation in a model where multiple parameters must be estimated is of interest. This is because, as mentioned in Chapter 3, current parameter estimation methods generally rely on the parameter space being low dimensional so as to avoid issues in tracking highly non-linear covariances (Santitissadeekorn & Jones 2015). As FTLE-ABC does not require covariance tracking the algorithm may not suffer from such issues and could offer potential advantages over current methods. The gridding method used to calculate the FTLE here may in this case become too computationally expensive, but using a more computationally efficient method such as those discussed in Section 2.5 may still allow FTLE-ABC to provide an advantage when the parameter dimension is high (Fortina et al. 2015).

More generally, implementing the algorithm with other coherent structure extraction techniques may allow the algorithm to be retro-fitted to the geophysical flow it is being applied to. For example, the Stochastic Sensitivity method mentioned in Section 2.2 (Balasuriya 2020b) is designed for systems with uncertainty in Lagrangian tracer trajectories and hence an ABC algorithm for parameter estimation based on this method may allow performance advantages when using data sets with high uncertainty in tracer positions. Similarly, another method mentioned in Section 2.2, the Lagrangian Averaged Vorticity Deviation (Haller et al. 2016), is designed for extracting rotational coherent structures and hence may afford an advantage for parameter estimation in systems dominated by rotational flow. If the algorithm is sufficiently generalised, the LCS method used within it could be selected to optimally exploit structural information present in the geophysical system of interest.

Drawing from the current literature in LaDA techniques, most successful techniques use a combination of methods within the same DA scheme, an approach which may also allow FTLE-ABC better performance. Slivinski et al. (2015) for example, find using an

EnKF to update high dimensional state variables and a PF to update lower dimensional Lagrangian drifters yields better results than a single method for updating drifters and states. Similarly, FTLE-ABC and potential variants may be more successful as the Lagrangian update component of a hybrid LaDA scheme.

There is a rich literature of computational methods, theory and applications of Lagrangian coherent structures from which the Lagrangian data assimilation literature could benefit. Few studies have rigorously investigated theoretical links between Lagrangian data assimilation methods, Lagrangian coherent structure theory and how interconnections between these two fields could be used to enhance the applicability and accuracy of current LaDA methods. A stronger theoretical basis for Lagrangian coherent data assimilation, such as by formulating the observation operator, forecast and observation covariance matrices in the language of LCS, could help guide the development of new LaCoDA methods. In Chapter 4, for example, all LaCoDA algorithms are formulated so that the coherent structure is calculated from one observation time to the next, to preserve the usual data assimilation assumption of independence between observations. However, coherent structures have their own time scales (c.f. Figure 6.1) that may be much longer than the time between observations. A potential question which arises then is how to set up the data assimilation problem if the observed coherent structures may have been calculated using observations from previous time steps. Similarly, how could the classic DA setup be altered if there are fine-scale coherent structures on a short timescale, and large-scale coherent structures on a longer timescale to exploit both kinds of structure? Such higher-level amalgamations of Lagrangian coherent structure theory and DA will benefit greatly from a strong theoretical basis connecting the two areas.

Although much is still to be investigated, the methods considered here are promising early developments in using LCS theory to improve LaDA techniques. This thesis realises one of many potential amalgamations of LCS and DA numerical methods and suggests such interdisciplinary fusions may continue to progress research which tackles chaos in geophysical modeling.

Bibliography

- Alligood, K., Sauer, T. & Yorke, J. (1997), *Chaos*, Springer, pp. 1–19.
- Apte, A. & Jones, C. K. R. T. (2013), ‘The impact of nonlinearity in Lagrangian data assimilation’, *Nonlinear Proc. Geophys.* **20**, 329–341.
- Apte, A., Jones, C. & Stuart, A. (2008), ‘A Bayesian approach to Lagrangian data assimilation’, *Tellus A: Dynamic Meteorology and Oceanography* **60**(2), 336–347.
URL: <https://doi.org/10.1111/j.1600-0870.2007.00295.x>
- Baines, P. G. (2008), ‘Lorenz, E.N. 1963: Deterministic non-periodic flow. Journal of the Atmospheric Sciences 20, 130–41.1’, *Progress in physical geography* **32**(4), 475–480.
- Balasuriya, S. (2016), ‘Meridional and Zonal Wavenumber Dependence in Tracer Flux in Rossby Waves’, *Fluids* **1**(3), 30.
- Balasuriya, S. (2020a), ‘Stochastic sensitivity: a computable lagrangian uncertainty measure for unsteady flows’, *SIAM Review* **62**, 781–816.
- Balasuriya, S. (2020b), ‘Uncertainty in finite time Lyapunov exponent calculations’, *Journal of Computational Dynamics* **7**(2), 313–337.
- Balasuriya, S., Kalampattel, R. & Ouellette, N. T. (2016), ‘Hyperbolic neighbourhoods as organizers of finite-time exponential stretching’, *Journal of Fluid Mechanics* **807**, 509–545.
- Balasuriya, S., Ouellette, N. T. & Rypina, I. I. (2018), ‘Generalized Lagrangian coherent structures’, *Physica D: Nonlinear Phenomena* **372**, 31–51.
URL: <https://doi.org/10.1016/j.physd.2018.01.011>
- Banicki, M. & Wiggins, S. (2010), ‘Finite-time Lagrangian transport analysis: stable and unstable manifolds of hyperbolic trajectories and finite-time Lyapunov exponents’, *Nonlinear Process. Geophys.* **17**, 1–36.
- Barreira, L. (2017), Introduction, in ‘Lyapunov Exponents’, Springer International Publishing, pp. 1–28.

- Bellsky, T., Berwald, J. & Lewis, M. (2013), ‘Non-global parameter estimation using local ensemble kalman filtering’, *arXiv.org: 10.1175/MWR-D-13-00200.1* .
- Beron-Vera, F., Olascoaga, M., Brown, M., Kocak, H. & Rypina, I. (2010), ‘Invariant-tori-like Lagrangian coherent structures in geophysical flows’, *Chaos* **20**(1).
- Bozorgmagham, A. E., Ross, S. D. & Schmale, D. G. (2013), ‘Real-time prediction of atmospheric Lagrangian coherent structures based on forecast data: An application and error analysis’, *Physica D: Non-linear Phenomena* **258**, 47–60.
- Brunton, S. & Rowley, C. (2010), ‘Fast computation of finite-time Lyapunov exponent fields for unsteady flows’, *Chaos* **20**, 017503.
- Carrassi, A., Bocquet, M., Bertino, L. & Evensen, G. (2018), ‘Data assimilation in the geosciences: An overview of methods, issues, and perspectives’, *Wiley interdisciplinary reviews. Climate change* **9**(5).
- Cencini, M. & Vulpiani, A. (2013), ‘Finite size Lyapunov exponent: review on applications’, *Journal of Physics A: Mathematical and Theoretical* **46**(25), 254019.
- Charney, J. (1951), *Dynamic Forecasting by Meteorological Process*, American Meteorological Society.
- Doerner, R., Hubinger, B., Martienssen, W., Grossmann, S. & Thomae, S. (1999), ‘Stable manifolds and predictability of dynamical systems’, *Chaos Solitons Fractals* **10**, 1759–1782.
- Evensen, G. (2009), *The Ensemble Kalman Filter*, 2 edn, Springer.
- Fortina, A., Briffard, T. & Garon, A. (2015), ‘A more efficient anisotropic mesh adaptation for the computation of Lagrangian coherent structures’, *Journal of Computational Physics* **285**, 100–110.
- Fowler, A. & McGuinness, M. (2019), Introduction, *in* ‘Chaos’, Springer, pp. 1–19.
- Froyland, G., Stuart, R. & van Sebille, E. (2014), ‘How well-connected is the surface of the global ocean?’, *Chaos* **24**, 033126.
- Gaultier, L., Djath, B., Verron, J., Brankart, J., Brasseur, P. & Melet, A. (2014), ‘Inversion of submesoscale patterns from a high-resolution Solomon Sea model: Feasibility assessment’, *Journal of Geophysical Research: Oceans* .
- Gaultier, L., Verron, J., Brankart, J.-M., Titaud, O. & Brasseur, P. (2013), ‘On the inversion of submesoscale tracer fields to estimate the surface ocean circulation’, *Journal of Marine Systems* **126**, 33–42.

- Gelman, A., Carlin, J. B., Stern, H. S., Dunson, D. B., Vehtari, A. & Rubin, D. B. (2013), 3 edn, Taylor & Francis Ltd.
- Gleick, J. (2008), The butterfly effect, *in* ‘Chaos: Making a new science’, Penguin.
- Haller, G. (2001), ‘Distinguished material surfaces and coherent structures in three-dimensional fluid flows’, *Physica D* **149**, 248–277.
- Haller, G. (2002), ‘Lagrangian coherent structures from approximate velocity data’, *Phys. Fluids A* **14**, 1851–1861.
- Haller, G. (2011), ‘A variational theory of hyperbolic Lagrangian coherent structures’, *Physica D* **240**, 574–598.
- Haller, G. (2015), ‘Lagrangian Coherent Structures’, *Annual Review of Fluid Mechanics* **47**(1), 137–162.
- Haller, G. & Beron-Vera, F. (2012), ‘Geodesic theory of transport barriers in two-dimensional flows’, *Physica D* **241**, 1680–16702.
- Haller, G., Hadjighasem, A., Farazmand, M. & Huhn, F. (2016), ‘Defining coherent vortices objectively from the vorticity’, *J. Fluid Mech.* **795**, 136–173.
- Haller, G. & Yuan, G. (2000), ‘Lagrangian coherent structures and mixing in two-dimensional turbulence’, *Physica D* **147**, 353–370.
- Husic, B. E., Schlueter-Kuck, K. L. & Dabiri, J. O. (2019), ‘Simultaneous coherent structure coloring facilitates interpretable clustering of scientific data by amplifying dissimilarity’, *J. Fluid Mech.* **861**, 886–900.
- Ide, K., Kuznetsov, L. & Jones, C. K. R. T. (2002), ‘Lagrangian data assimilation for point vortex systems’, *J. Turbulence* **3**.
- IPCC (2020), ‘Reports’.
URL: <https://www.ipcc.ch/reports/>
- Kalnay, E. (2002), Data assimilation, *in* ‘Atmospheric Modeling, Data Assimilation and Predictability’, Cambridge University Press, p. 136–204.
- Karrasch, D. (2012), ‘Comment on “A variational theory of hyperbolic Lagrangian coherent structures . [physica d 240 (2011) 574–598].”’, *Physica D* **241**, 1470–1473.
- Karrasch, D. & Haller, G. (2013), ‘Do finite-size Lyapunov exponents detect coherent structures?’, *Chaos* **23**, 043126.

- Katzfuss, M., Stroud, J. R. & Wikle, C. K. (2016), ‘Understanding the ensemble Kalman filter’, *The American Statistician: General* **70**(4), 350–357.
- Kew, E., Rossi, V., Sudre, J., Weimerskitch, H., Lopez, C., Hernandez-Garcia, E., Marsac, F. & Garcon, V. (2009), ‘Top marine predators track Lagrangian coherent structures’, *Proceedings of the National Academy of Science* **109**(106), 8245–8250.
- Koyama, H. & Watanabe, M. (2010), ‘Reducing forecast errors due to model imperfections using ensemble kalman filtering.’, *Mon. Wea. Rev.* **138**, 3316–3332.
- Kuhn, A., Rössl, C., Weinkauff, T. & Theisel, H. (2012), A benchmark for evaluating FTLE computations, in ‘2012 IEEE Pacific Visualization Symposium’, pp. 121–128.
- Kunsch, H. R. (2013), ‘Particle filters’, *Bernoulli* **19**(4), 1391–1403.
URL: <http://www.jstor.org/stable/23525756>
- Kuznetsov, L., Ide, K. & Jones, C. (2003), ‘A method for assimilation of Lagrangian data’, *The Monthly Weather Review* **131**(10), 2247—2260.
- Lahoz, W. A. & Schneider, P. (2014), ‘Data Assimilation: making sense of Earth Observation’, *Frontiers in Environmental Science* **2**, 16.
- Law, K., Stuart, A. & Zygalakis, K. (2015), *Data Assimilation: Mathematical Introduction*, Springer-Verlag.
- Leeuwen, P. (2015), Nonlinear data assimilation, in ‘Frontiers in Applied Dynamical Systems: Reviews and Tutorials’, Vol. 2, Springer, chapter Nonlinear Data Assimilation for high-dimensional systems.
- Leeuwen, P., Kunsch, H., Nerger, L., Potthast, R. & Reich, S. (2019), ‘Particle filters for high-dimensional geoscience applications:a review’, *Q J R Meteorol Soc.* **145**, 2335–2365.
- Lekien, F., Coulliette, C., Haller, G., Paduan, J. & Marsden, J. (2007), ‘Optimal pollution release in Monterey Bay based on non-linear analysis of coastal radar data’, *Environmental Science and Technology* **41**(18), 6562–6572.
URL: <https://doi.org/10.1021/es0630691>
- Lekien, F., Coulliette, C., Mariano, A., Ryan, E., Shay, L., Haller, G. & Marsden, J. (2005), ‘Pollution release tied to invariant manifolds: A case study for the coast of Florida’, *Physica D* **210**(1), 1–20.
URL: <https://doi.org/10.1016/j.physd.2005.06.023>
- Lekien, F. & Ross, S. D. (2010), ‘The computation of finite-time Lyapunov exponents on unstructured meshes and for non-Euclidean manifolds’, *Chaos* **20**, 017505.

- Lin, M., Xu, M. & Fu, X. (2017), ‘GPU-accelerated computing for Lagrangian coherent structures of multi-body gravitational regimes’, *Astrophysics and Space Science* **362**(4), 1–11.
- Lipinski, D. & Mohseni, K. (2010), ‘A ridge tracking algorithm and error estimate for efficient computation of Lagrangian coherent structures’, *Chaos* **20**, 017504.
- Lorenz, E. (1963), ‘Deterministic nonperiodic flow’, *Journal of Atmospheric Science* **20**(2), 130–141.
- Lu, F., Weitzel, N. & Monahan, A. H. (2019), ‘Joint state-parameter estimation of a nonlinear stochastic energy balance model from sparse noisy data’.
URL: <http://arxiv.org/abs/1904.05310>
- Lyapunov, A. (1892), The general problem of the stability of motion, PhD thesis, The University of Moscow.
- Macleán, J., Santitissadeekorn, N. & Jones, C. (2017), ‘A coherent structure approach for parameter estimation in Lagrangian Data Assimilation’, *Physica D: Nonlinear Phenomena* **360**, 36–45.
- Madja, A. & Harlim, J. (2012), The search for efficient particle filters for high-dimensional turbulent dynamical systems, in ‘Filtering Complex Turbulent Systems’, Cambridge University Press, pp. 316–347.
- May-Newman, K., Vu, V. & Herold, B. (2016), ‘Modeling the link between left ventricular flow and thromboembolic risk using Lagrangian coherent structures’, *Fluids* **1**(4).
- McIntyre, M. & Palmer, T. (1983), ‘Breaking Planetary Waves in the Stratosphere’, *Nature* **305**(5935), 593–600.
- Mezić, I., Loire, S., Fonoberov, V. A. & Hogan, P. (2010), ‘A new mixing diagnostic and gulf oil spill movement’, *Science* **310**, 486–489.
- Mihailovic, D. T., Mimi, G. & Arsenic, I. (2014), ‘Climate predictions: The chaos and complexity in climate models’, *Advances in Meteorology* **19**, 1–2014.
- Mirona, P., Vétela, J., Garona, A., Delfourba, M. & El Hassana, M. (2015), ‘Anisotropic mesh adaptation on Lagrangian coherent structures’, *Journal of Computational Physics* **231**, 6419–6437.
- NASA Image of the day Gallery* (2020).
- Nolte, D. D. (2019), *Nonlinear Dynamics and Chaos*, 2 edn, Oxford University Press.

- Norgard, G. & Bremer, P. (2012), ‘Second derivative ridges are straight lines and the implications for computing Lagrangian coherent structures’, *Physica D.* **241**, 1475–1476.
- Olascoaga, M., Beron-Vera, F., Brand, L. & Koçak, H. (2008), ‘Tracing the early development of harmful algal blooms on the West Florida Shelf with the aid of Lagrangian coherent structures’, *Journal of Geophysical Research: Oceans* **113**(C12).
- Olçay, A., Pottebaum, T. & Krueger, P. (2010), ‘Sensitivity of Lagrangian coherent structure identification to flow field resolution and random errors’, *Chaos* **20**, 017506.
- Onu, K., Huhn, F. & Haller, G. (2014), ‘LCS tool : A computational platform for Lagrangian coherent structures’.
URL: <http://proxy.library.adelaide.edu.au/login?url=https://search-proquest-com.proxy.library.adelaide.edu.au/docview/2078499976?accountid=8203>
- Peacock, T. & Dabiri, J. (2010), ‘Introduction to focus issue: Lagrangian coherent structures’, *Chaos* **20**, 017501.
- Petoukhov, V., Rahmstorf, S., Petri, S. & Schellnhuber, H. J. (2013), ‘Quasi-resonant amplification of planetary waves and recent Northern Hemisphere weather extremes’, *Proceedings of the National Academy of Sciences* **110**(14), 5336–5341.
URL: <https://doi.org/10.1073/pnas.1222000110>
- Pierrehumbert, R. (1991a), ‘Chaotic mixing of tracer and vorticity by modulated travelling rossby waves’, *Geophysical & Astrophysical Dynamics* **58**(1-4), 285–319.
- Pierrehumbert, R. (1991b), ‘Large-scale horizontal mixing in planetary atmospheres’, *Physics of Fluids A: Fluid Dynamics* **3**(5), 1250–1260.
URL: <https://doi.org/10.1063/1.858053>
- Pierrehumbert, R. T. & Yang, H. (1993), ‘Global chaotic mixing on isentropic surfaces’, *J. Atmos. Sci.* **50**(15), 2462–2480.
- Reich, S. & Cotter, C. (2015), Basic data assimilation algorithms, in ‘Probabilistic Forecasting and Bayesian Data Assimilation’, Cambridge University Press, pp. 316–347.
- Ross, S., Tanaka, M. & Senatore, C. (2010), ‘Detecting dynamical boundaries from kinematic data in biomechanics’, *Chaos* **20**, 017507.
- Salman, H., Ide, K. & Jones, C. K. R. T. (2008), ‘Using flow geometry for drifter deployment in Lagrangian data assimilation’, *Tellus A: Dynamic Meteorology and Oceanography* **60**(2), 321–335.

- Salman, H., Kuznetsov, L., Jones, C. K. R. T. & Ide, K. (2006), ‘A method for assimilating Lagrangian data into a shallow-water equation ocean model’, *Mon. Weather. Rev.* **134**, 1081–1101.
- Samelson, R. & Wiggins, S. (2006), *Steadily Translating Waves and Meanders*, Springer.
- Santitissadeekorn, N. & Jones, C. K. R. T. (2015), ‘Two-Stage Filtering for Joint State-Parameter Estimation’, *Mon. Weather Rev.* **143**, 2028–2042.
- Schlueter-Kuck, K. L. & Dabiri, J. O. (2019), ‘Model parameter estimation using coherent structure colouring’, pp. 1–21.
- Scranton, K., Knappe, J. & de Valpine, P. (2014), ‘An approximate Bayesian computation approach to parameter estimation in a stochastic stage-structured population model’, *Ecology* **95**(5).
URL: <https://doi-org.proxy.library.adelaide.edu.au/10.1890/13-1065.1>
- Shadden, S. (2007), ‘Lagrangian coherent structures in n-dimensional systems’, *Journal of Mathematics and Physics* **48**, 65404.
- Shadden, S. C., Lekien, F. & Marsden, J. E. (2005), ‘Structure and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows’, *Physica D: Nonlinear Phenomena* **212**(3), 271–304.
URL: <https://doi.org/10.1016/j.physd.2005.10.007>
- Sisson, S. A., Fan, Y. & Tanaka, M. M. (2007), ‘Sequential monte carlo without likelihoods’, *Proceedings of the National Academy of Sciences* **104**(6), 1760–1765.
- Sivakumar, B. (2004), ‘Chaos theory in geophysics: past, present and future’, *Chaos, Solitons and Fractals* **19**(2), 441–462.
- Slivinski, L., Spiller, E., Apte, A. & Sandstede, B. (2015), ‘A Hybrid Particle-Ensemble Kalman Filter for Lagrangian Data Assimilation’, *Monthly Weather Review* **143**(1), 195–211.
- Snyder, C., Bengtsson, T., Bickel, P. & Anderson, J. (2008), ‘Obstacles to High-Dimensional Particle Filtering’, *Monthly Weather Review* **136**(12), 4629–4640.
URL: <http://proxy.library.adelaide.edu.au/login?url=https://search-proquest-com.proxy.library.adelaide.edu.au/docview/198074608?accountid=8203>
- Solonen, A., Ollinaho, P., Laine, M., Haario, H., Tamminen, J. & Jarvinen, H. (2012), ‘Efficient mcmc for climate model parameter estimation: Parallel adaptive chains and early rejection.’, *Bayesian Anal.* **7**(3), 715–736.

- Spiller, E., Budhiraja, A. & Jone, C. K. R. T. (2008), ‘Modified particle filter methods for assimilating Lagrangian data into a point-vortex model’, *Physica D* **237**, 1498–1506.
- Stewart, I. (1997), Strange attractors, in ‘Does god play dice? The new mathematics of chaos’, Penguin UK.
- Surace, S. C., Kutschireiter, A. & Pfister, J. (2019), ‘How to Avoid the Curse of Dimensionality: Scalability of Particle Filters with and without Importance Weights’, *SIAM Review* **61**(1), 79–91.
URL: <https://doi.org/10.1137/17M1125340>
- Szanyi, S., Lukovich, J. V. & Barber, D. (2016), ‘Lagrangian analysis of sea-ice dynamics in the Arctic Ocean’, *Polar Research* **35**(1).
- Tang, W., Chan, P. & Haller, G. (2010), ‘Accurate extraction of Lagrangian coherent structures over finite domains with application to flight data analysis over Hong Kong International Airport’, *Chaos* **10**(1), 017502–1–8.
- Tina, T., Welch, D., Strelkova, N., Ipsen, A. & Stumpf, M. P. H. (2008), ‘Approximate Bayesian Computation scheme for parameter inference and model selection in dynamical systems’, *J. R. Soc. Interface* **6**, 187–202.
- Urban, N. & Keller, K. (2010), ‘Probabilistic hindcasts and projections of the coupled climate, carbon cycle and atlantic meridional overturning circulation system: a bayesian fusion of century-scale observations with a simple model.’, *Tellus A* **62**(4), 737–750.
- Vetra-Carvalho, S., van Leeuwen, P. J., Nerger, L., Barth, Alexander and Altaf, M. U., Brasseur, P., Kirchgessner, P. & Becker, J. (2018), ‘State-of-the-art stochastic data assimilation methods for high-dimensional non-Gaussian problems’, *Tellus A: Dynamic Meteorology and Oceanography* **70**, 1–43.
URL: <https://doi.org/10.1080/16000870.2018.1445364>
- Villagran, A., Huerta, G., Jackson, C. S. & Sen, M. K. (2008), ‘Computational Methods for Parameter Estimation in Climate Models’, *Bayesian Analysis* **3**(4), 823—850.
- von Hardenberg, J., Fraedrich, K., Lunkeit, F. & Provenzale, A. (2000), ‘Transient chaotic mixing during a baroclinic life cycle’, *Chaos* **10**, 122–134.
- Wang, Y., Gozolchiani, A., Ashkenazy, Y., Berezin, Y., Guez, O. & Havlin, S. (2013), ‘Dominant Imprint of Rossby Waves in the Climate Network’, *Physical Review Letters* **111**(13), 138501(1)–138501(5).